

NRAO-CHTC HTCondor collaboration

Created by James Robnett [X], last modified by K. Scott Rowe 48 minutes ago

- Current Questions
- In progress
- Answered Questions

Current Questions

tokens and collector.locate

It seems that if the submit host is HTC23, you need a user token in order for the API (nraosync specifically) to locate the schedd.

```
import os
import classad
import htcondor

def upload_file():
    try:
        ads = classad.parseAds(open(os.environ['_CONDOR_JOB_AD'], 'r'))
        for ad in ads:
            try:
                globaljobid = str(ad['GlobalJobId'])
                print("DEBUG: globaljobid is", globaljobid)
            except:
                return(-1)
        except Exception:
            return(-1)

        print("DEBUG: upload_file(): step 1in")
        submithost = globaljobid.split("#")[0]
        print("DEBUG: submithost is", submithost)
        collector = htcondor.Collector()
        print("DEBUG: collector is", collector)
        schedd_ad = collector.locate(htcondor.DaemonTypes.Schedd, submithost)
        print("DEBUG: schedd_ad is ", schedd_ad)

        upload_file()
```

This code works if both the AP and EP are version 10. But if the AP is version 23 then it fails weather the EP is version 10 or version 23. It works with version 23 iff I have a ~/condor/tokens.dnmpost token. Why do I need a user token to run collector.locate against a schedd?

I was going to test this on CHTC but I can't seem get an interactive job on CHTC anymore.

transfer\_output\_files change in version 23

My silly nraosync transfer plugin relies on the user setting transfer\_output\_files = job.ad in the submit description file to trigger the transfer of files. Then my nraosync plugin takes over and looks at nrao\_output\_files for the files to copy. But with version 23, this no longer works. I am guessing someone decided that internal files like .job.ad, .machine.ad, \_condor\_stdout, and \_condor\_stderr will no longer be tranferable via transfer\_output\_files. Is that right? If so, I think I can work around it. Just wanted to know.

getenv

Did it change since 10.0? Can we still use getenv in DAGs or regular jobs?

ikrowe Nov 5 2024: getenv no longer includes your entire environment as of version 10.7 or so. But instead it only includes the environment variables you list with the "ENV GET" syntax in the .dag file.

https://git.ligo.org/groups/computing/-/epics/30

Installing version 23

I am looking at upgrading from version 10 to 23 LTS. I noticed that y'all have a repo RPM to install condor but it installs the Feature Release only. It doesn't provide repos to install the LTS.

https://htcondor.readthedocs.io/en/main/getting-htcondor/from-our-repositories.html

Virtual memory vs RSS

Looks like condor is reporting RSS but that may actually be virtual memory. At least according to Felipe's tests.

ANSWER: Access to the cgroup information on the nmpost cluster is good because condor is running as root and condor reports the RSS accurately. But on systems using glidein like PATH and OSG they may not have appropriate access to the cgroup so memory reporting on these clusters may be different than memory reporting on the nmpost cluster. On glide-in jobs condor reports the virtual memory across all the processes in the job.

CPU usage

Felipe has had jobs put on hold for too much cpu usage.

```
runResidualCycle_n4.imycle8.condor.log.012 (269680.000.000) 2024-07-18 17:17:03 Job was held.
runResidualCycle_n4.imycle8.condor.log- Excessive CPU usage. Please verify that the code is configured to use a limited number of cpus/threads, and matches request_cpus.
```

GREG: Perhaps only some machines in the OSPool have checks for this and may be doing something wrong or strange.

2024-09-16: Felipe asked about this again.

Missing batch\_name

A DAG job, submitted with hundreds of others, doesn't show a batch name in condor\_q, just DAG: 371239. Just one job, all the others submitted from the same template do show batch names

Ausre/aaocluster/pipeline/Vlass\_prod/spool/quicklook/VLASS3.2\_T1727.J201445+263000\_P172318v1\_2024\_07\_12T16\_40\_09.270

```
nmpost-master krowe >condor_q -dag -name mcitroy -g -all
--
vlapipe vlass_q1.dag+370186 7/16 10:30 1 1 _ _ 3 370193.0
vlapipe vlass_q1.dag+370191 7/16 10:31 1 1 _ _ 3 370194.0
vlapipe DAG: 371239 7/16 10:56 1 1 _ _ 3 371536.0
--
```

GREG: Probably a condor bug. Try submitting it again to see if the name is missing again.

WORKAROUND: condor\_qedit job.id JobBatchName "asdflsd"

DAG failed to submit

Another DAG job that was submitted along with hundreds of others looks to have created vlass\_q1.dag.condor.sub but never actually submitted the job. condor.log is empty.

Ausre/aaocluster/pipeline/Vlass\_prod/spool/quicklook/VLASS3.2\_T1813.J093830+283000\_P175122v1\_2024\_07\_06T16\_33\_34.742

ANSWERS: Perhaps the schedd was too busy to respond. Need more resources in the workflow container?

Need to handle error codes from condor\_submit\_dag. 0 good. 1 bad. (chausman)

Setup /usr/bin/mail on mcitroy so that it works. Condor will use this to send mail to root when it encounters an error. Need to submit jira ticket to SSA. (krowe)

Resubmitting Jobs

I have an example in

Ausre/aaocluster/pipeline/Vlass\_prod/spool/se\_continuum\_imaging/VLASS2.1\_T1030.J194602-033000\_P161384v1\_2020\_08\_15T01\_21\_14.433

of a job that failed on nmpost106 but then HTCondor resubmitted the job on nmpost105. The problem is the job actually did finish, just got an error transferring back all the files, so when the job was resubmitted, it copied over an almost complete run of CASA which sort of makes a mess of things. I would rather HTCondor just fail and not re-submit the job. How can I do that?

```
022 (167287.000.000) 2023-12-24 02:43:57 Job disconnected, attempting to reconnect
Socket between submit and execute hosts closed unexpectedly
Trying to reconnect to slot1_1@nmpost106.aoc.nrao.edu <10.64.2.180-9618&alias=nmpost106.aoc.nrao.edu&noUDP&sock=stardt_5795_776a>
...
023 (167287.000.000) 2023-12-24 02:43:57 Job reconnected to slot1_1@nmpost106.aoc.nrao.edu
stardt address: <10.64.2.180-96187&adrs=10.64.2.180-9618&alias=nmpost106.aoc.nrao.edu&noUDP&sock=stardt_5795_776a>
stardt address: <10.64.2.180-96187&adrs=10.64.2.180-9618&alias=nmpost106.aoc.nrao.edu&noUDP&sock=slot1_1_39813_8c2c_400>
...
007 (167287.000.000) 2023-12-24 02:43:57 Shadow exception!
Error from slot1_1@nmpost106.aoc.nrao.edu: Repeated attempts to transfer output failed for unknown reasons
0 - Run Bytes Sent By Job
0 - Run Bytes Received By Job
...
040 (167287.000.000) 2023-12-24 02:45:09 Stared transferring input files
Transferring to host: <10.64.2.178-96187&adrs=10.64.2.178-9618&alias=nmpost105.aoc.nrao.edu&noUDP&sock=slot1_13_163338_25ab_452>
...
040 (167287.000.000) 2023-12-24 03:09:22 Finished transferring input files
...
001 (167287.000.000) 2023-12-24 03:09:22 Job executing on host: <10.64.2.178-96187&adrs=10.64.2.178-9618&alias=nmpost105.aoc.nrao.edu&noUDP&sock=stardt_5724_c431>
```

ANSWER: Maybe



## Singularity at PATH

My singularity jobs run but get the following error output

```
INFO Discarding path 'hadoop'. File does not exist
INFO Discarding path 'ceph'. File does not exist
INFO Discarding path 'hdfs'. File does not exist
INFO Discarding path 'hizarf'. File does not exist
INFO Discarding path 'imrthadoop'. File does not exist
INFO Discarding path 'imrthdfs'. File does not exist
WARNING: Environment variable HAS_SINGULARITY already has value [True], will not forward new value [] from parent process environment
WARNING: Environment variable REQUIRED_OS already has value [default], will not forward new value [] from parent process environment
/srv/gwms-user-job-wrapper.sh: line 882: /usr/bin/singularity: No such file or directory
```

WORKAROUND:

```
universe = vanilla
+SingularityImage = "path/to/myimage"
or
+SingularityImage = "docker://debian"
```

2023-02-06 krowe: I sent mail to Christina about this and she also suggested the workaround. She said they are working on a proper fix.

## Blocking on upload

Don't have condor block on the transfer plugin uploading. It doesn't block on download. When it blocks on upload and the upload is large, the job may get killed if NOT\_RESPONDING\_TIMEOUT isn't set to something larger than the 3600 second default.

## stdout and stderr with plugins

When using a transfer plugin to transfer output files, stdout and stderr are copied back as \_condor\_stdout and \_condor\_stderr. It doesn't rename them to what output and error are set to in the submit description file. If I use a transfer plugin for just input files and not output files, then stdout and stderr are copied back as requested in the submit description file.

This seems like a bug to me. Since my plugin isn't transferring these files, that means HTCondor is doing it so HTCondor should honor what is set in the submit description file whether I am using a transfer plugin or not.

Perhaps have rsync transfer these instead? Or use another custom classad instead of output\_destination? Or what if I put \_condor\_stdout in +nrao\_output\_files?

Actually, it doesn't seem to be triggered by just output\_destination. If I set output\_destination = \$ENV(PWD) and don't use a plugin for output files, I get stdout.jobid.log like I requested.

From [coatsworth@cs.wisc.edu](mailto:coatsworth@cs.wisc.edu) Mon Nov 29 12:26:12 2021

```
I've looked into this in the file transfer code. On the execution
side, we always write stdout and stderr to the _condor_stdout and
_condor_stderr files, then we remap them back to user-provided names
after a job completes. When you have output_destination set, our File
Transfer mechanism does not send files back to the submit machine by
default. However since your plugin is explicitly rsync-ing files back
there, they get moved without going through the remapping.
```

I think your File Transfer Mechanism **does** send files back to the submit machine by default. My transfer plugin is not transferring \_condor\_stderr nor \_condor\_stdout.

Apr. 22, 2022 krowe: Actually I think that since I set 'output\_destination = nraosync://.' in the submit description file, the FTM "let" using the plugin. It has to because output\_destination requires that everything use the plugin. So the FTM calls the plugin with \_condor\_stdout and \_condor\_stderr which activates the upload\_file() function in the plugin and the files are copied. This is why they aren't remapped. I configured the plugin to also copy these files and rename them. Perhaps instead I could watch for them in upload\_file() and remap them there?

ANSWER: Greg is going to tell Mark to put all this plugin work on the back burner or maybe stop altogether. Our plugin works with its work-arounds so this stuff is not critical.

## nraosync\_plugin.py

Since HTCondor walks the directories in transfer\_output\_files and submits the files one at a time the plugin, which doesn't work with rsync, we decided to work around the problem.

```
# Trick HTCondor into launching the plugin to handle output files
transfer_output_files = .job.ad
# custom job ad of files/dirs using nraosync_plugin.py
+nrao_output_files = "software data"
```

Call our upload\_rsync() function before calling upload\_file() in main()

```
with open(args['outfile'], 'w') as outfile:
    # krowe Oct 28 2021:
    # if args['upload']:
    #     if upload_rsync() != 0:
    #         raise err
    for ad in infile_ads:
```

ANSWER: I explained this to CHTC. They think it is at least an elegant hack. :~)

## Transfer Plugin Upload

Working with Mark Coatsworth on this.

I have added my [nraosync\\_plugin.py](#) to /usr/libexec/condor on the execution host and added the following configuration to the execution host:

```
FILETRANSFER_PLUGINS = $(LIBEXEC)/nraosync_plugin.py, $(FILETRANSFER_PLUGINS)
```

I have the following job:

```
#/bin/sh
mkdir newdir
date > newdir/date
/bin/sleep ${1}
```

and the following submit file:

```
executable = smaller.sh
arguments = "27"
output = stdout.$(ClusterId).log
error = stderr.$(ClusterId).log
log = condor.$(ClusterId).log
should_transfer_files = YES
transfer_input_files = /Users/krowe/.ssh/condor_transfer
transfer_output_files = newdir
output_destination = nraosync://$ENV(PWD)
+WantOPProxy = True
queue
```

The resulting input file that is fed to my plugin when the plugin is called with the -upload argument ([nraosync\\_plugin.in](#)) contains this:

```
[ LocalFileName = "/lustre/aoc/admin/tmp/condor/testpost003/execute/dir_29453/_condor_stderr"; Uri = "nraosync://lustre/aoc/sciops/krowe/plugin/_condor_stdout" ] [ LocalFileName = "/lustre/aoc/admin/tmp/condor/testpost003/execute/dir_29453/_condor_stdout"; Uri = "nraosync://lustre/aoc/sciops/krowe/plugin/_condor_stdout" ]
[ LocalFileName = "/lustre/aoc/admin/tmp/condor/testpost003/execute/dir_29453/newdir/date"; Uri = "nraosync://lustre/aoc/sciops/krowe/plugin/newdir/date" ]
```

I am surprised to see that it sets LocalFileName and Uri to the file inside newdir instead of newdir itself. Needless to say, this makes rsync unhappy as newdir doesn't exist on the destination yet.

If I create 'newdir' in the destination directory before submitting the job, the plugin will correctly copy the 'date' file back to the 'newdir' directory but the condor log file shows the following:

```
022 (4149.000.000) 08/05 09:22:04 Job disconnected, attempting to reconnect
Socket between submit and execute hosts closed unexpectedly
Trying to reconnect to slot1_4@testpost003.aoc.nrao.edu <10.64.1.173:9618?address=10.64.1.173:9618&alias=testpost003.aoc.nrao.edu&noUDP&socks=stand_28565_cae3>
...
023 (4149.000.000) 08/05 09:22:04 Job reconnected to slot1_4@testpost003.aoc.nrao.edu
stand address: <10.64.1.173:9618?address=10.64.1.173:9618&alias=testpost003.aoc.nrao.edu&noUDP&socks=stand_28565_cae3>
starter address: <10.64.1.173:9618?address=10.64.1.173:9618&alias=testpost003.aoc.nrao.edu&noUDP&socks=slot1_4_28601_bde4_d612>
...
```

condor re-runs the upload portion of the plugin four more times before finally giving up with this error

```
007 (4149.000.000) 08/05 09:22:31 Shadow exception!
Error from slot1_4@testpost003.aoc.nrao.edu: Repeated attempts to transfer output failed for unknown reasons
0 - Run Bytes Sent By Job
1007 - Run Bytes Received By Job
```

If I create a file like 'outfile' instead of 'newdir' and transfer that, everything works fine.

I have an example in [home/mu\\_kscott/htcondor/plugin\\_small](#)

ANSWER: Greg will look into this. K. Scott is working with Mark Coatsworth on this.

## condor\_off vs condor\_drain

I would like to be able to issue a command to an execute host telling it to stop accepting new jobs and let the current jobs finish. I would also like that host to stay in the condor\_status output with a message indicating what I have done (i.e. draining, offline, etc) I think I want something that does some of condor\_off and some of condor\_drain. Is there such a beast?

For example, a -peaceful option to condor\_drain might be perfect.

condor\_off

- Pro: Immediately prevents new jobs from starting on the node with the -startd option
- Pro: Is supposed to let existing jobs finish running with the -peaceful option but
- Con: Stopped working with -peaceful for me sometime after upgrading to 9.0.4
- Con: Doesn't have a -reason option

```

    • Con: removes host from condor_status
condor_drain

    • Pro: Leaves host in condor_status
    • Pro: Has a -reason option
    • Con: Immediately evicts all running jobs on node

ANSWER: condor_status -master and use condor_off

ANSWER: Greg thinks condor_drain should have a -peaceful option. (bug)

2022-10-05 krowe: HTCondor 9.12.0 "Added -drain option to condor_off and condor_restart". I think this might be the solution I wanted they just went in a different direction. Instead of 'condor_drain -peaceful' there is now a 'condor_off -drain'. The feature isn't in the LTS release yet. Perhaps it will be in 9.0.18. Then I will test it.

2023-07-20 krowe: condor_off now (as of some version before 10.0.2) has a -reason option. You can see the -reason and -drain options with --help but they aren't in the man page.

2023-07-20 krowe: I tried condor_off with HTCondor-10.0.2. It now has the -drain option which seems to kill jobs. Not what I was expecting. Also, the -reason option can only be used with the -drain option. When I use condor_on to put the node back into the cluster, it restarted my job. Again, I think I really want a -peaceful argument to condor_drain. That would be the best solution for me.
```

### Show offline nodes

Say I set a few nodes to offline with a command like `condor_off -startd -peaceful -name nmpost120` How can I later check to see which nodes are offline?

- `condor_status -offline` returns nothing
- `condor_status -long nmpost120` returns nothing about being offline
- The following shows nodes where startd has actually stopped but it doesn't show nodes that are set offline but still running jobs (e.g. Retiring)
  - `condor_status -master -constraint "STARTD_StartTime == 0"`
- This shows nodes that are set offline but still running jobs (a.k.a. Retiring)
  - `condor_status |grep Retiring`

ANSWER: 2022-06-27

condor\_status -const Activity == "Retiring"

offline ads, which is a way for HTCondor to update the status of a node after the startd has exited.

condor\_drain -peaceful # CHTC is working on this. I think this might be the best solution.

### Glidein

The only documentation I can find on glidein (<https://htcondor.readthedocs.io/en/latest/grid-computing/introduction-grid-computing.html#highlight-glideinintroduction>) seems to imply that glidein only works with Globus "HTCondor permits the temporary addition of a Globus-controlled resource to a local pool. This is called glidein." Is this correct? Is there better documentation? Is glidein even a technology or software package or is it just a generic term?

ANSWER: Greg will look at re-writing this.

### request\_virtualmemory

If I set `request_virtualmemory = 2G`, `condor_submit` accepts it as a valid knob but the job stays idle and never runs.

```

    request_memory = 1G
    request_virtualmemory = 2G
```

If I set `request_virtualmemory = 2000000`, which should be the same as 2G, the job runs but doesn't set `memory.memsw.limit_in_bytes` in the cgroup.

Oct. 11, 2021 krowe: Checked with HTCondor-9.0.6. Problem still exists unchanged.

ANSWER: krowe sent mail to Greg about it

## Answered Questions

- JOB ID question from Daniel
  - When I submit a job, I get a job ID back. My plan is to hold onto that job ID permanently for tracking. We have had issues in the past with Torque/Maui because the job IDs got recycled later and our internal bookkeeping got mixed up. So my questions are:
    - Are job IDs guaranteed to be unique in HTCondor?
    - How unique are they—are they `_globally_` unique or just unique within a particular namespace (such as our cluster or the submit node)?
    - A Job ID (ClusterID.ProcID)
    - DNS name of the schedd and ctime of the job\_queue.log file.
    - It is unique to a schedd.
    - We should talk with Daniel about this. They should craft their own ID. It could be seeded with a JobID but should not depend on just it.
  - Upgrading HTCondor without killing jobs?
    - schedd can be upgraded and restarted without losing state assuming the restart is less than the timeout.
    - currently restarting execute services will kill jobs. CHTC is working on improving this.
    - negotiator and collector can be restarted without killing jobs.
    - CHTC works hard to ensure 8.8.x is compatible with 8.8.y or 8.9.x is compatible with 8.9.y.
  - Leaving data on execution host between jobs (data reuse)
    - Todd is working on this now.
    - Ask about installation of CASA locally and ancillary data (dcache)
      - CHTC has a Ceph filesystem that is available to many of their execution hosts (notably the larger ones)
      - There is another software filesystem where CASA could live that is more used for admin usage but might be available to us.
      - We could download the tarball each time over HTTP. CHTC uses a proxy server so it would often be cached.
    - Environment: Is there a way to have condor "login" when a job starts thus sourcing `etc/profile` and the user's `rc` files? Currently, not even \$HOME is set.
    - A good analogy is Torque does a `su - $username` while HTCondor just does a `su $username`.
    - WORKAROUND: setting `getenv = True` which is like the `-V` option to `qsub`, may help. It doesn't source `rc` files but does inherit your current environment. This may be a problem if your current environment is not what you want on the cluster node. Perhaps the cluster node is a different OS or architecture.
    - ANSWER: condor doesn't execute things with a shell. You could set your executable as `bash` and then have the arguments be the executable you used to have. I just changed our stuff to statically set \$HOME and I think that is good enough.
  - Flocking: Suppose I have two hosts in the same pool. testpost-master is a submit-host and testpost-serv-1 is both a submit-host and the central-manager. testpost-serv-1 is configured to flock to CHTC but testpost-master is not. Is it possible to submit a job on testpost-master that will flock to CHTC by somehow leveraging testpost-serv-1? In other words, do I have to setup flocking and an external IP on every submit host?
    - ANSWER: There isn't a good way to do this. So eventually we will need to make testpost-master flock to CHTC and possibly remove the ability of testpost-serv-1 to flock.
  - It seems the transfer mechanism won't transfer symlinks to directories (e.g. `data/vlass.ms -- /lustre/aoc/...`) Is there a way around this?
    - ANSWER: there is no flag to chase symlinks at the moment. The top level dir (e.g. `data`) could be a symlink if transfer\_input\_files=data but it will then transfer the contents of data instead of data itself.
    - If symlink -- data and transfer\_input\_files=symlink I get the error Transfer of symlink to directories is not supported.
    - If symlink -- /data and transfer\_input\_files=symlink I transfer the contents not the directory. In other words I don't have a data directory in scratch I have a VCLASS... directory.
    - If data/VCLASS -- /somepath/VCLASS and transfer\_input\_files=data/VCLASS/
    - If data/VCLASS -- /somepath/VCLASS and transfer\_input\_files=data/ I get the error Transfer of symlinks to directories is not supported.
  - DAG log time stamps, is there a way to differentiate data import/export time and process run time.
    - Look in the job log file not the dag log file
    - 040 (150.000.000) 2020-06-15 13:05:45 Started transferring input files
      - Transferring to host: <10.64.10.172:9618?address=10.64.10.172:9618&alias=nmpost072.aoc.nrao.edu&noUDP&sock=slot1\_1\_72656\_7984\_40>
    - ...
    - 040 (150.000.000) 2020-06-15 13:06:04 Finished transferring input files
  - Rank and Preemption: Can we use Rank to set "preferences" without requiring job preemption?
    - ANSWER: There are 2 kinds of rank (job rank, machine rank). job rank (RANKJob... in submit file) is purely a preference. That does not preempt. Machine rank (in startd.config) will preempt. NEGOTIATOR\_PRE\_JOB\_RANK is a third type of rank that works at a pool level and is often used to pack jobs efficiently.
  - Update on software store for CASA either on shared Ceph storage or admin software storage
    - Staging area for datasets 100MB - 1Tb. This is where we could try keeping the dcache assuming doing so doesn't overwhelm the filesystem.
      - /staging/nu\_jobnettt
      - Requirements = {Target,HasCHTCStaging == true}
      - Quota: 100GB, 100K files
    - Squid area for 100MB - 1GB input or shared software. This is where we could keep casa.tgz and then have the execution host retrieve it via HTTP.
      - /squid/nu\_jobnettt
      - only accessible via this path on the submit hosts. Execution hosts will need to access it via HTTP.
      - transfer\_input\_files = http://proxy.chtc.wisc.edu/SQID/nu\_jobnettt/casa.tgz
    - Software area We can use this in run-time applications. Think of it like `lsr/local`.
      - /software/nu\_jobnettt/casa/casa-pipeline-release-5.6.1-8.e17
      - export PATH=\$(optifocalbin)/software/nu\_jobnettt/casa/casa-pipeline-release-5.6.1-8.e17/bin:\$PATH
      - Quota: 5GB, 100K files
  - Public\_input\_files: How is this different than transfer\_input\_files and when would one want to use it instead of files or URLs with transfer\_input\_files?
    - <https://htcondor.readthedocs.io/en/latest/users-manual/file-transfer.html#public-input-files>
    - This is still a work in progress. It may allow for caching on a squid server, fetchable by others someday.
  - Flocking: When we flock to CHTC what is the data path for transfer\_input\_files? Is it our submit host and CHTC's execution host, is CHTC's submit host involved?
    - Dataflow is from our schedd (submit host) to their execute host but CCB will reverse the connection. Their execution hosts are publicly addressable but that may not be necessary.
  - How can we choose the data path for transfer\_input\_files to our clients given multiple networks. Currently we assume it will use the 1Gb link but we have IB links. Is there a way for condor to use the IB link just for transferring files, is that hostname based? Other ideas?
    - CHTC doesn't have a good solution for this.
    - We could upgrade from 1Gb to 10Gb
    - We could use the IB names for everything (problematic for submit hosts that don't have IB)
    - We could not use transfer mechanism and instead use something else like scp
    - We could use a custom transfer plugin
  - Are there known issues with distributed scratch via NFS or Lustre w.r.t tmpdir or other, e.g. OpenMPI complains about tmpdir being on network FS?
    - Some problems with log files on the submit host but rare.
  - Any general best practices to support MPI in terms of class ads or other.
    - Use the shared memory transport for security
  - Is there a way DAGMan can be told to ignore errors, in some cases we want a DAG to mindlessly continue vs retry.
    - The job is considered successful based on the return of the post script. If there isn't a post script, the success is based on the return of the job.
  - Transfer mechanism: Documentation implies that only files with an mtime newer than when the transfer\_input\_files finished will be transferred back to the submit host. While running a dag, the files in my working directory (which is in both transfer\_input\_files and transfer\_output\_files) seem to always have an mtime around the most recent step in the DAG suggesting that the entire working directory is copied from the execution host to the submit host at the end of each DAG step. Perhaps this means the transfer mechanism only looks at the mtime of the files/dirs specified in transfer\_output\_files and doesn't descend into the directories.
    - Subdirectories are treated differently
    - SOLUTION: I think casa just touches every file and therefore condor is forced to copy everything in the working directory. I have been unable to reproduce the problem outside of casa.
    - SOLUTION: If you specify a directory, HTCondor will transfer the entire directory not just files with new mtime.
  - Does the transfer mechanism accept any sort of regular expression? E.g. `transfer_input_files="*.txt"`
    - No
  - Can the transfer mechanism accept manifest files? E.g. a file that is a list of files?
    - Use include : <some file> in the submit script where <some file> contains the full transfer\_input\_files line
    - Use queue FILES from manifest Which defines the submit variable \$(FILES) which could be used in a transfer\_input\_files like: transfer\_input\_files = \$(FILES)
    - Perhaps a plugin
  - What other options are there than holding a job? I find myself not noticing, sometimes for hours, that a job is on hold. Is there a way to make jobs fail instead of getting held? I assume others will make this mistake like me.

- I see I can set `periodic_remove = (JobStatus == 5)` but HTCondor doesn't seem to think that is an error so if I have notification = Error I don't get any email.
  - Greg will look into adding a Hold option to notification
  - The HTCondor idea of held jobs is that you submitted a large DAG of jobs, one step is missing a file and you would like to put that file in place and continue the job instead of the whole DAG failing and having to be resubmitted. This makes sense but it would be nice to be notified when a job gets held.
  - Greg wrote "notification = error in the submit file is supposed to send email when the job is held by the system, but there's a bug now where it doesn't. I'll fix this."
- What limits are there to transfer\_input\_files? I would sometimes get **Failed to transfer files** when the number of files was around 10,000
  - ANSWER: There is a memory limit because of ClassAds but in general there isn't a defined limit.
  - Is there a way to generate the dag dot file without having to submit the job?
    - The `-no_submit` option doesn't create the dot file
    - Is adding **NOOP** to all the JOB commands the right thing to do? The DAG still gets submitted but then quickly ends.
    - ANSWER: You need to submit the DAG. NOOP is the current solution
  - Is there a way to start a dag at a given point? E.g. if there are 5 steps in the dag, can you start the job at step 3?
    - Is the answer again to add **NOOP** to the JOB commands you don't want to run?
    - ANSWER: Splices may work here but NOOP is certainly a tool to use.
  - I see at CHTC jobs now require a `request_disk` setting. How does one launch interactive jobs?
    - ANSWER: This is a bug.
  - For our initial tests, we want to flock jobs to CHTC that transfer about 60GB input and output. Eventually we will reduce this significantly but for now what can we do?
    - Globus or rsync to move data? If Globus, how to do so in an automated way (E.g. no password)?
    - ANSWER: Using the Transfer Mechanism from our submit server to CHTC's execution host is ok with CHTC because it doesn't interfere with their submit host. Outbound from their execute hosts is also allowed (scp, rsync, etc).
  - Use `-Longjob = true` attribute for 72+ hour jobs. This is a CHTC-specific knob.
  - How can I know if my job swapped?
      - ANSWER: CHTC nodes have no or minimal swap space.
  - Condor Annex processing in AWS. Is there support for spot market
      - ANSWER: Condor Annex does indeed support the spot market. It is a bit more work to set up because you don't say "give my X of Y", but "I'll pay d1 dollars for machines like X1 and d2 for machines like X2, etc."
  - What network mask should we use to allow ssh from CHTC into NRAO? Is it a class B or several class Cs?
      - ANSWER: The ip (v4 I) ranges for CHTC execute nodes are
        - 128.104.100.0/24
        - 128.104.55.0/24
        - 128.104.58.0/23
        - 128.105.244.0/23
  - Is there a ganglia server or some other monitor service at CHTC we can view?
      - ANSWER: We have a bunch of ganglia and grafana graphs for the system, but I think they are restricted to campus folks and tend to show system-wide utilization and problem
  - I have a machine with an externally-accessable, non-NATed address (146.88.10.48) and an internal, non-routable address (10.64.1.226). I want to install condor\_annex on this machine such that I can submit jobs to AWS from it. I don't necessarily need to submit jobs to local execute hosts from this machine. Should I make this machine a central manager, a submit host, both, or does it matter?
      - ANSWER: I think instances in AWS will need to contact both the schedd (submit host) and collector (central manager) from the internet using port 9618. So either both submit host and central manager need external connections and IPs with port 9618 open or combine them into one host with an external IP and port 9618 open to the Internet.
    - Last time I configured condor\_annex I was using an older version of condor (8.8.3 I think) and used a pool password for security. Now I am using 8.9.7. Is there a newer/better security method I should use?
      - ANSWER: The annex still primarily uses pool password, so let's stick with that for now.
  - How can I find out what hosts are available for given requirements (LongJob, memory, staging)
      - condor\_status -compact -constraint "HasChtcStaging==true" -constraint "DetectedMemory>500000" -constraint "CanRunLongJobs isnt Undefined"
      - Answer: yes this is correct but it doesn't show what other jobs are waiting on the same resources. Which is fine.
      - It looks to me like most hosts at CHTC are setup to run LongJobs. The following shows a small list of about 20 hosts so I assume all others can run long jobs. Is the correct?
        - condor\_status -compact -constraint "CanRunLongJobs is Undefined"
        - JongJobs is for something like 72 hours. So it might be best to not set it unless we really need it like step23.
        - Answer: yes this is correct but it doesn't show what other jobs are waiting on the same resources. Which is fine.
    - Is port 9618 needed for flocking or just for condor\_annex?
        - ANSWER: Greg thinks yes 9618 is needed for both flocking and condor\_annex.
      - Are there bugs in the condor.log output of a DAG node? For example, I have a condor.log file that clearly shows the job taking about three hours to run yet at the bottom lists user time of 13 hours and system time of 1 hour. [https://open-confluence.nrao.edu/download/attachments/40541486/step07\\_py\\_condor.log?api=v2](https://open-confluence.nrao.edu/download/attachments/40541486/step07_py_condor.log?api=v2)
        - And as for the cpu usage report, there could very well be a bug, but first, is your job multi-threaded or multi-process? If so, the cpu usage will be the aggregate across all cpu cores.
        - Yes they are all parallel jobs to some extent so I accept your answer for that job. But I have another job that took 21 hours of wallclock time and yet the condor.log shows 55 minutes of user and 5:34 hours of system time. [https://open-confluence.nrao.edu/download/attachments/40541486/step05\\_py\\_condor.log?api=v2](https://open-confluence.nrao.edu/download/attachments/40541486/step05_py_condor.log?api=v2)
        - ANSWER: if you look, the user time is actually 6 days and 55 minutes. I missed the 6 in there.
    - Given a DAG where some steps are to **always** run at AOC and some are to **always** run in CHTC how can we dictate this. Right now local jobs flock to CHTC if local resources are full. Can we make local jobs idle instead of flock?
        - ANSWER: Use PoolNames. I need to make a testpost PoolName.
        - PoolName = "TESTPOST"
        - STARTD\_ATTRS = \$(STARTD\_ATTRS) PoolName
    - It seems that when using DAGs the recommended method is to define variables in the DAG script instead of submit scripts. This makes sense as it allows for only one file, the DAG script, that needs to be edited to make changes. But, is there a way to get variables into the DAG script from the command line or environment or an include\_file or something?
        - ANSWER: There is an INCLUDE syntax but there is no command-line or environment variable way to get vars into a DAG.
      - We are starting to run 10s of jobs in CHTC requiring 40GB as part of a local DAG. Are there any options we can set to improve their execution chance. What memory footprint (32, 20, 16, 8GB) would significantly improve their chances.
        - ANSWER: only use +LongJob if the job needs more than 72 hour, which is the default "walltime".
      - How can we set AWS Tags with condor\_annex? We'd like this to track jobs and set billing tags. Looks like there isn't really a way.
        - SOLUTION: Greg wrote (Oct. 5, 2020) "we've just now added code to the condor\_annex to support setting multiple aws tags on the annex command line". K. Scott expects it will take a while before that code makes it to released software.
        - Launch Templates didn't work. I don't think condor\_annex supports Launch Templates.
        - Use aws-user-data options to condor\_annex?
          - I have tried all sorts of user-data and default-user-data file options. On-demand apparently no longer works and I was never able to get something working with spot-fleet. I think all things user-data are non-functional.
          - I tried setting a tag in the role defined in config.json (aws-ec2-spot-fleet-tagging-role) but that tag didn't translate to the instance.
          - I tried adding a tag to the AMI when creating a new AMI (EC2 -- Instances -- Actions -- Image -- Create Image). Didn't work.
          - What about self-tagging? The instance figures out its instance id and runs aws.
            - wget -qO- http://instance-data/latest/meta-data/instance-id
              - returns nothing when logged in as nobody (condor\_ssh\_to\_job)
              - returns nothing when logged in as centos (ssh -i ~/ssh/...)
              - returns instanceid when logged in as root (so as centos then sudo su)
              - Aha! There is a **firewall (ptables) rule blocking exactly this**. But I can't figure out what file sets this iptables rule on boot.
            - I tried adding tags to the json file using both ResourceType set to instance and spot-fleet-request. Neither created an instance with my tag.
      - What:
          - "TagSpecifications": [
            - {
            - "ResourceType": "instance",
            - "Tags": [
              - {
              - "Key": "budget",
              - "Value": "VLASS"
              - }
              - }
              - ]
              - }
              - ]
              - }
      - What are the clever solutions to submitting N different DAG jobs with each having different parameters?
          - T1034
            - J220200-003000
              - bin, working, data
            - J220600-003000
              - bin, working, data
            - ...
          - T1035
            - J170743-393000
              - bin, working, data
            - J171241-383000
              - bin, working, data
            - ...
          - ANSWERS:
            - INCLUDE syntax for DAGs
            - include syntax for submit files
            - make a template of files
            - use a PRE script that populates things
            - usedagdir
      - I had a job killed because it exceeded 72 hours even though I set `-LongJobs = true` in the submit file
          - 2385.0 krowe 9/22 20:43 Error from slot\_1\_1@e2008.chtc.wisc.edu: Job failed to complete in 72 hrs
          - ANSWER: the knob is singular `-LongJob = true`
        - What are the options to setting up HTCondor to both flock to CHTC and annex to AWS? Multiple submit hosts? Multiple CMs? etc.
          - ANSWER: Philosophy is for everyone to submit in one place and let condor sort out where it goes.
          - CHTC locks annex jobs to a different CM that actually starts the annex.
            - Submit annex job on SM. If then flocks to a different CM that can create the annex
        - As we feared, referencing the cache of convolution functions (cfcache) directly from staging performed poorly. This is due to a (fsat?) pathology that fares poorly on distributed filesystems. Jobs ran 3 to 4 times faster when we copied cfcache from /staging to local disk. I ran a small data set test with full parameters at CHTC that copied cfcache from /staging to local disk and step05 took only 16.7 hours instead of the 56.8 hours it had taken using cfcache on /staging.
      - Condor\_annex bug: Edit [fixurlibexec/condor/condor-annex-ec2](#) and comment out the line `chkconfig condor || exit 1` because this line is a hold-over from older versions that put condor in init.d. Now that it is in systemd, this line causes condor to exit.
        - SOLUTION: Greg submitted ticket on this.
        - Debugging held jobs. I had thought that **when to transfer\_output = ON\_EXIT\_OR\_EVICT** would copy the scratch area back to the submit machine so files there could be inspected. But that doesn't seem to happen for me.
          - <https://htcondor.readthedocs.io/en/latest/users-manual/file-transfer/renamingcopying-it-and-when-to-transfer-files>
          - Evict means your job was killed because of policy like priority or time but not memory or disk.
          - Hold is often because of an error like missing file transfer or out of memory
        - SOLUTION: there isn't a good solution to hold.
        - Memory issue: Greg did find a bug deep in the code that may cause jobs to be killed because of memory issues. HTCondor occasionally gets a short read when looking at the process table via /proc, and then something like 2/3 of the processes are missing.
          - SOLUTION: CHTC will work towards a solution.
      - How can we have the .dag\* files written to a different directory? -usedagdir doesn't help.
          - ANSWER: There isn't a way to tell condor\_submit\_dag where to put the logs
        - Is there a for-loop structure available to DAG scripts or a range mechanic?
          - No
        - If 8.9.9 requires Globus from EPEL then it may have trouble being installed on a Globus endpoint because the EPEL version of Globus conflicts with the [Globus.org](#) version.
          - I told them about it. I have not tried installing HTCondor-8.9.9 so I am only guessing it will be a problem.

```

• Is there a recommended way to start annexes from a DAG? We have been using PRE scripts but sometimes it seems to fail.
  ◦ SOLUTIONS:
    • CHTC is working on a BEGIN syntax (provision) that will block a DAG node from starting until the annex is ready.
    • We could have the script not return until the annex is ready.
    • We could also have the job require a specific name that the create_annex creates.

• How can I set a variable in a DAG file that I can then use in the submit file in a conditional? None of the following seem to work
  ◦ DAG:
    • VARS step01 CHTC=""
    • VARS step05 CHTC="True"
  ◦ Submit:
    • if defined $(CHTC)
      • requirements = PoolName == "CHTC"
    • endif
  ◦ or
  ◦ DAG:
    • #VARS step01 CHTC="True"
    • VARS step05 CHTC="True"
  ◦ Submit:
    • if defined $(CHTC)
      • requirements = PoolName == "CHTC"
    • endif
  ◦ or
  ◦ DAG:
    • VARS step01 CHTC="False"
    • VARS step05 CHTC="True"
  ◦ Submit:
    • chtc_var = $(CHTC)
    • if $(chtc_var)
      • requirements = PoolName == "CHTC"
    • endif
  ◦ even though when I pass $(chtc_var) as arguments to the shell script, the shell script sees it as True.
  ◦ or
  ◦ DAG:
    • VARS node1 file="chtc.htc"
    • VARS node2 file="aws.htc"
  ◦ Submit:
    • include : $(file)

10/20/20 08:54:36 From submit: ERROR: on Line 9 of submit file:
10/20/20 08:54:36 From submit: Submit:-1:Error "", Line 0, Include Depth 1: can't open file
10/20/20 08:54:36 From submit:
10/20/20 08:54:36 From submit: ERROR: Failed to parse command file (line 9).
10/20/20 08:54:36 failed while reading from pipe.
10/20/20 08:54:36 Read so far: Submitting job(s)ERROR: on Line 9 of submit file: Submit:-1:Error "", Line 0, Include Depth 1: can't open fileERROR: Failed to parse command file (line 9).
10/20/20 08:54:36 ERROR: submit attempt failed

  ◦ Yet I can use a variable defined in a DAG for things like arguments and request_memory.
  ◦ I can also use file = $CHOICE(myindex, chtc.htc, aws.htc) where myindex is defined in a DAG it will set $(file) to the file I want to include but again if I use include : $(file) I get an error

10/20/20 11:58:58 From submit: Submitting job(s)ERROR on Line 13 of submit file: $CHOICE() macro: myindex is invalid index!
10/20/20 11:58:58 failed while reading from pipe.
10/20/20 11:58:58 Read so far: Submitting job(s)ERROR on Line 13 of submit file: $CHOICE() macro: myindex is invalid index!
10/20/20 11:58:58 ERROR: submit attempt failed

• Perhaps use requirements. Greg will send an example
  ◦ SOLUTION:
  ◦ DAG:
    • JOB step05 step05.htc
    • #VARS step05 SITE="chtc"
    • #VARS step05 SITE="aws"
  ◦ Submit:
    • #NRAOAtr = $(SITE)
    • Requirements = MyNRAOAtr == "chtc" ? PoolName == "CHTC" : PoolName != "CHTC"
    • Requirements = MyNRAOAtr == "chtc" ? (Target.HasCHTCStaging == true) : (Target.HasCHTCStaging != true)
    • myannex = "krowe-annex"
    • #MayUseAWS = True
    • Requirements = MyNRAOAtr == "aws" ? AnnexName == $(myannex) : AnnexName != $(myannex)
  ◦ I would set myannex in the DAG but when I do that it tries to find an AnnexName of "krowe - annex" (note spaces)

• ANSWER: My conclusion is that there are limitations on what one can do with variables in the submit file that were defined in the DAG file.

• Is there a config option that will cause condor to not start? We have diskless nodes and it is easier to modify the config file then change systemd.
  ◦ SOLUTION: Either set START_MASTER = False or START_DAEMONS = False depending on desired outcome.

• Torque has this command called pbsnodes that can not only offline/drain a node but keeps a note about it that all can see in one place. I know I can use condor_off to drain a node but is there a central place keep notes so I can remember a month later why I set a certain node to drain?
  ◦ ANSWER: there is no place to keep such notes but Greg likes the idea and may look into it.
  ◦ May want to use condor_drain instead of condor_off. condor_off will kill the startd when all jobs finish and it no longer shows up in condor_status. condor_drain will leave the node in condor_status.
  ◦ condor_drain doesn't work for me because it immediately sets jobs idle instead of letting them run to completion. This is why I use condor_off_startd-peaceful instead.

• How can you tell which job is associated with an email given the email message doesn't include a working dir or the assigned batch_name?
  ◦ CHTC will look into adding such information to the email condor sends.

• Bug in condor_annex: Underscores in the AnnexName prevent the annex from moving into the pool.
  ◦ Also when I try to terminate an annex with underscores (e.g. krowe_annex_casa5) with the command condor_off-annex krowe_annex_casa5 I get the following error
    • Found no ClassAds when querying pool (local)
    • Can't find addresses for master's for constraint 'AnnexName =?= "krowe_annex_casa5"'
    • Perhaps you need to query another pool.
  ◦ Greg has noted this bug

• Bug in condor_annex: The following will wait for an annex named krowe - annex - casa5 (note the spaces). If I pass $(myannex) as an argument to a shell script, the spaces are not there.
  ◦ include .htc
  ◦ submit .htc
    • myannex = krowe-annex-casa5
  ◦ submit .htc
    • include : include.htc
    • executable = /bin/sleep
    • arguments = 127
    • #MayUseAWS = True
    • requirements = AnnexName == $(myannex)
  ◦ queue
  ◦ Actually, I think this isn't a bug but a limitation on using macros. The AnnexName needs to be quoted but how can I quote a macro? Note, I have the same problems with AnnexNames that don't have hyphens (E.g. krowetest).
    • No: requirements = AnnexName == $(myannex)
    • No: myannex = "krowe-annex-casa5"
    • No: myannex = "krowe-annex-casa5"
    • No: myannex = "krowe-annex-casa5"
    • No: myannex = "krowe-annex-casa5"
  ◦ Idea: #annex = "krowe-annex-casa5"
  ◦ requirements = AnnexName == my.annex
  ◦ Greg has noted this bug

```

```

• Torque has this command called pbsnodes that can not only offline/drain a node but keeps a note about it that all can see in one place. I know I can use condor_off to drain a node but is there a central place keep notes so I can remember a month later why I set a certain node to drain?
  ◦ ANSWER: there is no place to keep such notes but Greg likes the idea and may look into it.
  ◦ May want to use condor_drain instead of condor_off. condor_off will kill the startd when all jobs finish and it no longer shows up in condor_status. condor_drain will leave the node in condor_status.
  ◦ condor_drain doesn't work for me because it immediately sets jobs idle instead of letting them run to completion. This is why I use condor_off_startd-peaceful instead.

```

```

• How can you tell which job is associated with an email given the email message doesn't include a working dir or the assigned batch_name?
  ◦ CHTC will look into adding such information to the email condor sends.

```

```

• Bug in condor_annex: Underscores in the AnnexName prevent the annex from moving into the pool.
  ◦ Also when I try to terminate an annex with underscores (e.g. krowe_annex_casa5) with the command condor_off-annex krowe_annex_casa5 I get the following error
    • Found no ClassAds when querying pool (local)
    • Can't find addresses for master's for constraint 'AnnexName =?= "krowe_annex_casa5"'
    • Perhaps you need to query another pool.
  ◦ Greg has noted this bug

• Bug in condor_annex: The following will wait for an annex named krowe - annex - casa5 (note the spaces). If I pass $(myannex) as an argument to a shell script, the spaces are not there.
  ◦ include .htc
  ◦ submit .htc
    • myannex = krowe-annex-casa5
  ◦ submit .htc
    • include : include.htc
    • executable = /bin/sleep
    • arguments = 127
    • #MayUseAWS = True
    • requirements = AnnexName == $(myannex)
  ◦ queue
  ◦ Actually, I think this isn't a bug but a limitation on using macros. The AnnexName needs to be quoted but how can I quote a macro? Note, I have the same problems with AnnexNames that don't have hyphens (E.g. krowetest).
    • No: requirements = AnnexName == $(myannex)
    • No: myannex = "krowe-annex-casa5"
    • No: myannex = "krowe-annex-casa5"
    • No: myannex = "krowe-annex-casa5"
    • No: myannex = "krowe-annex-casa5"
  ◦ Idea: #annex = "krowe-annex-casa5"
  ◦ requirements = AnnexName == my.annex
  ◦ Greg has noted this bug

```

## Nodesfree

How can one see nodes that are entirely unclaimed?

SOLUTION: condor\_status -const PartitionableSlot && Cpus == TotalCpus

## HERA queue

I want a proper subset of machines to be for the HERA project. These machines will only run HERA jobs and HERA jobs will only run on these machines. This seems to work but is there a better way?

machine config	submit file
HERA = True	requirements = (HERA == True)
STARTD_ATTRS = \$(STARTD_ATTRS) HERA	+partition = "HERA"
START = \$(STARTD) && (TARGET.partition !=?= "HERA")	

SOLUTION: yes, this is good. Submit Transforms could also be set on herapost-master (Submit Host)

<https://hecondor.readthedocs.io/en/latest/misc-concepts/transforms.html#howto-submit%50transform>

## Reservations

What if you know certain nodes will be unavailable for a window of time say the second week of next month. Is there a way to schedule that in advance in HTCondor? For example in Slurm

**scntrol create reservation starttime=2021-02-07T08:00:00 duration=7-0-0 nodes=nmpos[020-030] user=root reservationname=siv2022**

ANSWER: HTCondor doesn't have a feature like this.

## Bug: All on one core

```

• Bug where James's jobs are all put on the same core. Here is top-u krowe showing the Last Used Cpu (SMP) after I submitted five sleep jobs to the same host.
  ◦ Is this just a side effect of condor using cpusack instead of cpuset in cgroup?
  ◦ Is this a failure of the Linux kernel to schedule things on separate cores?
  ◦ Is this because cpu.shares is set to 100 instead of 1024?
  ◦ Check if CPU affinity is set in /proc/self/status
  ◦ Is sleep cpu-intensive enough to properly test this? Perhaps submit a while 1 loop instead?

```

```

PID USER PR NI VIRT RES SHR S %CPU WHEN TIME+ COMMAND P
66713 krowe 20 0 4364 348 288 S 0.0 0.0 0:00.01 sleep 22
66714 krowe 20 0 4364 352 288 S 0.0 0.0 0:00.02 sleep 24
66715 krowe 20 0 4364 348 288 S 0.0 0.0 0:00.01 sleep 24
66719 krowe 20 0 4364 348 288 S 0.0 0.0 0:00.02 sleep 2
66722 krowe 20 0 4364 352 288 S 0.0 0.0 0:00.02 sleep 22

```



```

GSI5003:Failed to authenticate. Globus is reporting error (851968:50). There
is probably a problem with your credentials. (Did you run grid-proxy-init?)
AUTHENTICATE:1004:Failed to authenticate using KERBEROS
AUTHENTICATE:1004:Failed to authenticate using FS

```

ANSWER:

condor\_submit-remote does indeed, tell the condor\_submit tool to submit to a remote schedd. (it also implies -spool)

Because the schedd can run the job as the submitting owner, and runs the shadow as the submitting owner, the remote schedd needs to not just authorize the remote user to submit jobs, but must authenticate the remote user as some allowed user.

Condor's IP host-based authentication is just authentication, it can say "everyone coming from this IP address is allowed to do X, but I don't know who that entity is".

So, for remote submit to work, we need some kind of authentication method as well, like IDTOKENS, munge.

## Authentication

- We're currently using host based authentication. Is there a future proof recommended authentication system for HTCondor-9.x for a site planning to use both on-premise cluster and CHTC flocking and or glide-ins to other facilities? host\_based? password? Tokens? SSL? Munge? Munge might be my preferred method as Slurm already requires it.
- If we're using containers for submit hosts is there a preferred authentication scheme (host based doesn't scale well).
  - ANSWER: idtokens

## HTCondor+Slurm

- Do people do HTCondor glide-ins to slurm where the HTCondor jobs are not preempted, as a way to share resources with both schedulers?
  - ANSWER: You can glide in to Slurm.
  - You can have Slurm preempt HTCondor jobs in favor of its own jobs (HTCondor jobs presumably will be resubmitted)
  - You can have HTCondor preempt Slurm jobs in the same sort of way.

## Transfer Plugin Order

HTCondor guarantees that the condor file transfer happens before the plugin transfer, but only when using the "multi-file" plugin style, like we have in our curl plugin. If you used the curl plugin as the model for rsync, you should be good.

## AMQP

The AMQP gateway that we had developed was called Qpid, and worked by tailing the user job log and turning it into qpid events. I suspect there's also ways to have condor plugins directly send amqp events as well.

## CPU Shares

Torque uses cpusets which is pretty straight forward, but HTCondor uses cpu.shares which confuses me a bit. For example, a job with **request\_cpus = 8** executing on a 24-core machine gets **cpu.shares = 800**. If there are no other jobs on node, does this job essentially get more CPU time than 1024/800?

ANSWER: yes it is opportunistic. If there are no other jobs running on a node you essentially get all the node.

## Nodescheduler

We found a way to implement our nodescheduler script in Slurm using the --exclude option. Is there a way to exclude certain hosts from a job? Or perhaps a constraint that prevents a job from running on a node that is already running a job of that user? Is there a better way than this?

```
requirements = Machine != "nmpost097.aoc.nrao.edu" && Machine != "nmpost119.aoc.nrao.edu"
```

```
badmachines=one+two+three
```

```
requirements not in $(badmachines)
```

I didn't get the actual syntax from Greg and I am apparently not able to look it up. The long syntax I suggested should work I just don't know what Greg's more efficient syntax is.

## condor\_ssh\_to\_job

Is there a way to use condor\_ssh\_to\_job to connect to a job submitted from a different submit host (schedd) or do you have to run it from the submit host used to submit the job? I have tried using the -name option to condor\_ssh\_to\_job but I always get **Failed to send GET\_JOB\_CONNECT\_INFO to schedd**

ANSWER: idtokens. Host-based and poassword are not sufficient to identify users and allow for this (and probably condor\_submit-remote).

## HTCondor Workshop vs Condor Week

ANSWER: Essentially it is "Condor Week Europe". Mostly the same talks but different customer presentations. Could be interesting for the different customer presentations.

## Shutdown

```

STARTD.DAEMON_SHUTDOWN = State == "Unclaimed" && Activity == "Idle" && (MyCurrentTime - EnteredCurrentActivity) > 600
MASTER.DAEMON_SHUTDOWN = STARTD_StartTime == 0

```

But I was running a job when it shut down.

```
07/19/21 11:45:01 The DaemonShutdown expression "State == "Unclaimed" && Activity == "Idle" && (MyCurrentTime - EnteredCurrentActivity) > 600" evaluated to TRUE: starting graceful shutdown
```

Could this be because we use dynamic slots?

```

testpost-cm-vml krowe >condor_status
Name OpSys Arch State Activity LoadAv Mem
slot1@testpost001.aoc.nrao.edu LINUX X86_64 Unclaimed Idle 0.000 193
slot1_1@testpost001.aoc.nrao.edu LINUX X86_64 Claimed Busy 0.000
slot1@testpost002.aoc.nrao.edu LINUX X86_64 Unclaimed Idle 0.000 144
slot1_1@testpost002.aoc.nrao.edu LINUX X86_64 Claimed Busy 0.810 49
slot1@testpost003.aoc.nrao.edu LINUX X86_64 Unclaimed Idle 0.000 193

```

I see that with dynamic slots, the parent slot (slot1) seems always unclaimed and idle and the child slots (slot1\_1) are Claimed and Busy. So I tried checking the ChildState attribute which looks to be a list but doesn't behave like one. For example, none of these show any slots

```

condor_status -const 'ChildState == ("Claimed")'
condor_status -const 'sum(ChildState) == 0'

```

Even though this produces true

```
classad_eval 'a = {}' 'sum(a) == 0'
```

ANSWER: Try this

```
condor_status -const 'size(ChildState) == 0'
```

## HTCondor and Slurm

NRAO has effectively two use cases: 1) Operations triggered jobs. These are well formulated pipeline jobs, they're still fairly monolithic and long running (many hours to few days). 2) User triggered jobs, these are of course not well formulated. We will be moving the operations jobs to htcondor. We plan to move the user triggered jobs to SLURM form Torque. There's enough noise in the two job loads that we don't want to have strict host carve outs for type 1 and type 2 jobs. What we anticipate doing is having a set of nodes known only to htcondor for the bulk of operations and a set of hosts controlled by SLURM for the user facing jobs. Periodically when they have a large set of operations jobs we'd like for them to burst into the SLURM controlled nodes. We neither anticipate nor want the slurm jobs to burst into the htcondor set of nodes.

Say we have two clusters (HTCondor and Slurm) and both can be submitted to from the same host. We want the HTCondor jobs to use the Slurm cluster resources when the HTCondor cluster resources are full, but we probably don't want to support preemption. How could we have HTCondor submit jobs to a Slurm cluster? (HTCondor-C, flocking, overlapping, batch-grid-type, HTCondor-CE, etc)

ANSWER: write our own factory that watched HTCondor and when it is full submit Pilot jobs to Slurm that launch startd daemons thus allowing the Payload jobs waiting in HTCondor to run. Will want to set the startd to exit after being idle for a little while, run the Pilot job as root, and figure out how to do cgroups properly.

## Shadow jobs and Lustre

We had some jobs get restarted because they lost contact with their shadow jobs. I assume this is because the shadow jobs keep the condor.log file open and if that file is on Lustre and Lustre goes down then the shadow job fails to communicate with the job and the job gets killed. Does that seem accurate to you?

```

nmpost-master root>ps auxww|grep shadow|grep krowe
krowe 1631810 0.0 0.38708 3676 ? S 09:29 0:00 condor_shadow -f 486.0 --schedd=<10.64.10.100:9618?addr=10.64.10.100-9618&noUDP&sock=5837_96cc_3> --xfer-queue=limit-upload,download,addr=<10.64.10.100:14115> <10.64.10.100:14115> -
nmpost-master root>ls -la /proc/1631810/fd
total 0
dr-x----- 2 root root 0 Jul 27 09:29 .
dr-xr-xr-x 8 krowe nmstaff 0 Jul 27 09:29 .
lr-x----- 1 root root 64 Jul 27 09:29 0 -> pipe:[16358528]
lr-x----- 1 root root 64 Jul 27 09:29 1 -> pipe:[16358540]
lwx----- 1 root root 64 Jul 27 09:29 18 -> socket:[16358529]
lwx----- 1 root root 64 Jul 27 09:29 2 -> pipe:[16358540]
lwx----- 1 root root 64 Jul 27 09:29 3 -> /lustre/aoc/sciops/krowe/condor.486.log
lwx----- 1 root root 64 Jul 27 09:29 4 -> socket:[16358542]

```

Here are some logs of a filed job

```

07/26/21 14:38:38 (479.0) (1188418): Job 479.0 is being evicted from slot1_1@nmpost114.aoc.nrao.edu
07/26/21 14:38:38 (479.0) (1188418): logEvictEvent with unknown reason (108), not logging.
07/26/21 14:38:38 (479.0) (1188418): **** condor_shadow (condor_SHADOW) pid 1188418 EXITING WITH STATUS 108

```

Exit Code 108 = can not connect to the condor\_startd or request refused

```
| 2021-07-26 14:16:39 (pid:91673) Lost connection to shadow, waiting 2400 secs for reconnect
```

ANSWER: Greg thinks this is an accurate description of the problem. Greg thinks this 2400 second timeout may be adjustable but do we want to? How long is long enough? Two choices: 1 decide we don't care 2 write log files to something other than Lustre.

## Rebooting Submit Host

What happens to running jobs if the submit host reboots? Shadow processes? What if the submit host is replaced with a new server? I think we have shown there is a 2400 second (40 minute) timeout.

ANSWER: state files are in \$(condor\_config\_val SPOOL) and you only have 40 minutes by default and that timeout is set at job submission time.

## Chirp in upload\_file

While I seem to be able to use chirp in the download\_file() function of a plugin, I cannot seem to use it in the upload\_file() portion. Something like the following will produce a line in the condor log file but not when executed from the upload\_file() function. This I have tested at CHTC.

```
| message = 'in upload_file()'
| subprocess.call(['usr/libexec/condor/condor_chirp', 'ulog', message])
```

I have an example in /home/nu\_kscott/htcondor/plugin\_small

The plugin hangs during the output and the processes running on the execution host look like this

```
| krowe 36107 0.0 0.0 58420 8128 ? Ss 13:44 0:00 condor_startd -f -local-name slot_type_1 -a slot_1_1 testpost-master.aoc.nrao.edu
| krowe 36571 0.8 0.0 182728 15004 ? Ss 13:46 0:00 /usr/bin/python3 /usr/libexec/condor/raorsync_plugin.py -rfile /lustre/aoc/admin/tmp/condor/testpost002/execute/dir_36107/raorsync_plugin.py.in -outfile /lustre/aoc/admin/tmp/condor/testpost002/execute/dir_36107/raorsync_plugin.py.out -upload
| krowe 36572 0.0 0.0 17084 1288 ? Ss 13:46 0:00 /usr/libexec/condor/condor_chirp ulog in upload_file()
```

If I kill the condor\_chirp process (36572), the plugin moves on to the next file to upload at which point it runs condor\_chirp again and hangs again. If I keep killing the condor\_chirp processes eventually the job finished properly.

ANSWER: Greg looked into this and said there is no good workaround. "This is simply a deadlock between chirp and the file transfer plugin. When transferring the output sandbox back to the submit machine, the HTCondor starter runs the file transfer code synchronously wrt the starter (it forks to do this while transferring the input sandbox...), and the starter also handles chirp calls."

## Timeout

What is the timeout setting called and can we increase it? Is it **JobLeaseDuration**? Can it be altered on a running job?

ANSWER: yes it is JobLeaseDuration and it can be changed in the execution host

## condor\_gpu\_discovery

I can't find the condor\_gpu\_discovery on my cluster (HTCondor-9.0.4) or CHTC (9.1.4) even on a GPU host.

ANSWER: /usr/libexec/condor/condor\_gpu\_discovery

## idtokens with RPMs

It seems that installing HTCondor-9.0.4 via RPMs doesn't automatically create signing key in /etc/condor/passwords.d/P00L like the documentation reads <https://htcondor.readthedocs.io/en/latest/admin-manual/security.html#highlight=idtokens&quick-configuration-of-security>

Also with the RPM install, **ALLOW\_WRITE = \*** which seems insecure. Does this even matter when **use security:recommended\_v9\_0**

ANSWER: this can probably just be ignored. Greg didn't think fresh installs actually created signing keys so this may be an error in documentation.

## idtokens

We are using HTCondor-9.0.4 and switched from using host\_based security to idtoken security with the following procedure.

On just the Central Manager named testpost-cm (which is the collector and schedd)

```
| openssl rand -base64 32 | condor_store_cred add -c -f /etc/condor/passwords.d/P00L
| condor token create -identity condor@testpost-cm.aoc.nrao.edu > /etc/condor/tokens.d/condor@testpost-cm.aoc.nrao.edu
| echo 'SEC_TOKEN_POOL_SIGNING_KEY_FILE = /etc/condor/passwords.d/P00L' >> /etc/condor/config.d/99-nrao
```

then switch to **use security:recommended\_v9\_0** in 00-htcondor-9.0.config

On the worker nodes (startds)

```
| scp testpost-cm:/etc/condor/passwords.d/P00L /etc/condor/passwords.d
| scp testpost-cm:/etc/condor/tokens.d/condor@testpost-cm.aoc.nrao.edu /etc/condor/tokens.d
| echo 'SEC_TOKEN_POOL_SIGNING_KEY_FILE = /etc/condor/passwords.d/P00L' >> /etc/condor/config.d/99-nrao
```

then switch to **use security:recommended\_v9\_0** in 00-htcondor-9.0.config

But then things like condor\_off don't work

```
| testpost-cm-vml root >condor_off -name testpost002
| ERROR
| AUTHENTICATE:1003:Failed to authenticate with any method
| AUTHENTICATE:1004:Failed to authenticate using SSL
| AUTHENTICATE:1004:Failed to authenticate using SCITOKENS
| AUTHENTICATE:1004:Failed to authenticate using GSI
| CSI:5002:Failed to authenticate. Globus is reporting error (851968:50). There is probably a problem with your credentials. (Did you run grid-proxy-init?)
| AUTHENTICATE:1004:Failed to authenticate using KERBEROS
| AUTHENTICATE:1004:Failed to authenticate using IDTOKENS
| AUTHENTICATE:1004:Failed to authenticate using F/S
| Can't send KIll-All-Daemons command to master testpost002.aoc.nrao.edu
```

ANSWER: The CONDOR\_HOST on the startds was not fully qualified. Also, both the startds and the collector/schedd were using the cname (testpost-cm) instead of the hostname (testpost-cm-vml). I changed them both to the following and now I can use both condor\_off and condor\_drain without error.

```
| CONDOR_HOST = testpost-cm-vml.aoc.nrao.edu
```

## Docs wrong for evaluating ClassAds?

This web page <https://htcondor.readthedocs.io/en/latest/man-pages/classads.html#highlight=evaluate&testing-classad-expressions> suggests that the following will produce **false** but for me it produces **error**

```
| condor_status -limit 1 -af 'regex( ".*tr*", "string" )'
```

ANSWER: The first asterisk shouldn't be there. This is a regex not globbing. Greg will look into updating this document.

Oct. 11, 2021 krowe: The documentation looks to have been corrected.

## Memory usage report

The memory usage report at the end of the condor log seems incorrect. I can watch the memory\_max\_usage\_in\_bytes in the cgroup get over 8,400MB yet the report in the condor log reads 6,464MB. Does the log only report the memory usage of the parent process and not include all the children? Is it an average memory usage over time?

ANSWER: It is a report of a sum of certain fields in memory.stat in the cgroup. Get Greg an example. Try it on two machines in case this is a problem of re-using the same cgroup. Or reboot and try again.

Oct. 11, 2021 krowe: With HTCondor-9.0.6, it looks like my tests are now reporting consistent values between memory\_max\_usage\_in\_bytes in the cgroup and Memory in the condor log. Except the memory\_max\_usage\_in\_bytes is in base-10 while the condor log is in base-2.

## Tracking jobs through various log files

What is the preferred method of tracking jobs through various log files like condor.log, StarterLog.slot1\_2, etc?

The condor.log uses a jobid but the StarterLogs use pid

ANSWER: condor.log to StarterLog on execute host to StarterLog.slot\* on execute host search for "Job <jobid>"

ANSWER: **condor\_history <jobid> -af LastRemoteHost** will give the slot id

## Flocking with idtokens

Does the following seem correct?

I setup rastan-vml as a standalone Central Manager, Schedd, and Startd (I'm starting to talk like an HTCondor admin now). This is what I had in the config on rastan-vml

```
| UID DOMAIN = aoc.nrao.edu
| JOB_DEFAULT_NOTIFICATION = ERRORCONDOR_ADMIN = krowe@nrao.edu
| CONDOR_HOST = rastan-vml.aoc.nrao.edu
| PoolName = "rastan"
| FLOCK_TO = testpost-cm-vml.aoc.nrao.edu
```

I then created a token for me in ~/condor/tokens.d but this did not allow jobs to flock from rastan-vml to testpost-cm.

I then copied the token from testpost-cm/etc/condor/tokens.d to rastan-vml/etc/condor/tokens.d and that was enough to get the job flocking.

ANSWER: Yes.

## gdrive example

I tried to use the gdrive plugin but couldn't find any documentation and failed to figure it out on my own.

ANSWER: ask coataworth

I swear this wasn't in the docs last week.

<https://htcondor.readthedocs.io/en/latest/users-manual/file-transfer.html?highlight=Transferring%20files%20to%20and%20from%20Google%20Cloud%20Storage#file-transfer-using-a-url>

But CHTC doesn't have a Google credential, so I can't use the gdrive plugin at CHTC.

Submitting job(s)  
OAuth error: Failed to securely read client secret for service gdrive; Tell your admin that gdrive\_CLIENT\_SECRET\_FILE is not configured

condor\_watch\_q

```
nmpost-master krowe >condor_watch_q
```

ERROR: Unhandled error: [Errno 2] No such file or directory: '/proc/sys/user/max\_inotify\_instances'. Re-run with -debug for a full stack trace.

ANSWER: it is in beta. send email to htcondor-admin about it

<https://htcondor.readthedocs.io/en/latest/overview/support-downloads-bug-reports.html>

Lauren suggest I ask for a ticket account.

Nevermind. It is slated to be fixed in 9.0.7.

## Launch numbers

Are there knobs to control how many jobs get launched at the same time and/or delay between launches? We are wondering because we hit our MaxStartups limit of 10:30:60 in sshd.

- FILE\_TRANSFER\_DISK\_LOAD\_THROTTLE ?
- JOB\_START\_COUNT /JOB\_START\_DELAY ?

ANSWER: JOB\_START\_COUNT looks like the right thing.

## plugin\_small

Can one of you please try the instructions in

/home/ru\_kscott/plugin\_small/small.htc

ANSWER: CHTC admins are required to use two-factor authentication via PAM. This means they can't use a passwordless ssh key in a job.

## Transferring back .ad files

I can add the following to my job and it will not cause an error but it also won't transfer the file

transfer\_output\_files = .job.ad, .machine.ad, .chrp.config

This isn't really important, I just thought it could be useful to diagnose jobs if I had a copy of the .job.ad and thought this would be a convenient way to get it. I am surprised that it neither causes and error, which it would if the file didn't actually exist, nor copies it. So I am guessing either the file is removed after it is checked for existence or HTCondor knows about its internal files and refuses to copy them.

ANSWER: Greg is not surprised by this.

## FILESYSTEM\_DOMAIN as requirement

I want to submit jobs that require a different filesystem but none of the following seem to work

requirements = (FILESYSTEM\_DOMAIN == "aoc.nrao.edu")  
FILESYSTEM\_DOMAIN = "aoc.nrao.edu"  
+FILESYSTEM\_DOMAIN = "aoc.nrao.edu"

Looks like the answer is

requirements = (FileSystemDomain == "aoc.nrao.edu")

<sarcasm>because that's perfectly obvious</sarcasm>

But let's say we have two clusters (aoc.nrao.edu and cv.nrao.edu) with different filesystems. I want jobs submitted in aoc.nrao.edu with a requirement of cv.nrao.edu to glidein to the cv.nrao.edu cluster. How can a factory script at cv.nrao.edu look for such jobs? I can't seem to use condor\_q -constraint to look for such jobs. The following doesn't work.

condor\_q -pool nmpost-cm-vml.aoc.nrao.edu -global -allusers -constraint 'Requirements = ((FileSystemDomain == "cv.nrao.edu"))'

ANSWER: I think the answer is not to use FileSystemDomain but to create our own custom classad like we do with the VLASS partition. Greg says it is possible to query for this requirement but the syntax is pretty gnarly. I think making a partition is a better solution.

## Removing tokens

Let's say I have a schedd that authenticates with an idtoken in /etc/condor/tokens.d. If I remove that token, I am still able to submit jobs from that host until condor is restarted. It has to be a restart as condor\_reconfig seems insufficient. This indicates to me that HTCondor is caching the token. Although it is strange that condor\_token\_list returns nothing immediately after removing the token yet HTCondor still can submit jobs. This is not really a problem but I was surprised by it and wanted to point it out in case it was unexpected. There doesn't seem to be a timeout either.

ANSWER: Greg knows about this. HTCondor establishes a relationship once authenticated and continues to use that relationship. It may timeout after 24 hours, not sure.

## Signing key

Given two separate clusters (testpost and nmpost), what should the signing keys and tokens look like?

Now that we use idtokens, I thought that to get a VM to be able to submit jobs I only needed to add our cluster's token to /etc/condor/tokens.d. But apparently I also need to add our cluster's signing key to /etc/condor/passwords.d. I since learned that this is probably because I created the signing key and token on our testpost cluster and then copied them to our nmpost cluster.

ANSWER: yes. create signing keys for each cluster.

## Jobs with a little swap

Say we had jobs that need 40GB of memory but occasionally, very briefly, spike to 60GB. With Torque this is not a problem because it will just let the job swap. It is not a big performance hit because the amount of time that memory is needed is very short compared to the runtime of the job. How could we do this in HTCondor? We really don't want to set a memory requirement of 60GB because we want to run multiple jobs on a node and doing so will significantly reduce the number of jobs we could put on a node.

Does the new DISABLE\_SWAP\_FOR\_JOB=false knob, introduced in 8.9.9, mean that HTCondor now swaps if needed by default?

ANSWER: try setting memory.swappiness for the condor cgroup.

ANSWER: The VLASS nodes don't have a swap partition. Make a swapfile on the vlass node (nmpost110) and see if that works.

## Allocated in the log file

If I submit a job at CHTC with request\_disk = 1 G the log output looks like

Partitionable Resources : Usage Request Allocated  
Cpus : 1 1  
Disk (KB) : 49 1048576 1485064  
IoHeavy : 0  
Memory (MB) : 1 1024 1024

But if I submit a job at CHTC with a request\_disk = 2 G the log output looks like

Partitionable Resources : Usage Request Allocated  
Cpus : 1 1  
Disk (KB) : 49 2097152 7258993  
IoHeavy : 0  
Memory (MB) : 0 1024 1024

What does the "Allocated" disk space mean in these examples?

ANSWER: with partitionable slots HTCondor allocates more disk space than you ask because then that slot might be used by a follow up job. This is because destroying and creating partitionable slots takes a full negotiation cycle which is measured in minutes.

MODIFY\_REQUEST\_EXPR\_REQUEST\_DISK=RequestDisk can alter this behavior. check docs. on the startd (execute host)

## Rebooting Execute Hosts

When an Execute Host unexpectedly reboots, what happens to the job? What are the options? Currently it looks like the job just "hangs". Condor\_q indicates that it is still running but it isn't. Looks like it eventually times out after the magic 40 minutes.

ANSWER: Correct

## condor\_off -reason

You added a -reason to condor\_drain, could the same be added to condor\_off?

ANSWER: Greg likes this idea and will look into it. Only recently did they implement offline ads that would allow this sort of thing.

ANSWER: daemons have a keep alive message. If startd expects keep alives from the starter if not received it gets killed. This is outside JobLeaseDuration. This is from the old days of Condor when it was scavaging cycles and didn't want to get in the user's way. Look into NOT\_RESPONDING\_TIMEOUT in the config file on the worker node. Default is

3600 seconds. Try setting it to something LARGE.

### Comments

This has probably already been mentioned but would it be possible to put comments after a condor command like so

```
batch_name = "test script" # dont show this
```

without the batch\_name being set to **test script # dont show this**

ANSWER: Not likely to be changed as doing so may break other things.

### Removing jobs with tokens

You can use tokens to remove jobs as other users but strangely not on the same host. For example: krowe and krowe2 have the same token (~/.condor/tokens.d/testpost). If I submit a job as krowe on testpost-master I cannot remove that job as krowe2 on testpost-master.

```
testpost-master$ condor_q -g -all -af clusterid owner jobstatus globaljobid
452 krowe 2 testpost-master.aoc.nrao.edu#452_0#1648820298
testpost-master$ condor_rm 452
Couldn't find/remove all jobs in cluster 452
testpost-master$ condor_rm -name testpost-master 452
testpost-master$ condor_rm -name testpost-master 452
Couldn't find/remove all jobs in cluster 452
```

However, if I submit a job as krowe on testpost-cm I "can" remove that job from testpost-master (condor\_rm -name testpost-cm 123). Is this a bug? Is it because when you are on the same host, HTCondor is trying UID authentication instead of token authentication? If so, is there a way to force to force token authentication?

ANSWER: Greg thinks this is because they choose the authentication type first and then stick with that type.

WORKAROUND: I "think" \_condor\_SEC\_DEFAULT\_AUTHENTICATION\_METHODS=IDTOKENS condor\_rm will use idtokens but Greg thinks this may not work so be warned.

### Condor Week

- transportation (airports, busses, cars, etc) Greg recommends a cab from the airport to campus.
- The Getting to the University of Wisconsin-Madison campus (By car, bus or plane) link under the Local Arrangements link is broken.
- The Ground Transportation from the Madison airport into download link seems to work but I bet you ment downtown.

### RADIAL CHTC support

- What role does CHTC have for this if any?
  - ANSWER: They PXE boot but use a local OS installation with puppet to keep them in sync.
- Singularity or Apptainer?
  - ANSWER: They are using Singularity now but may switch to Apptainer.

### Flocking and networking

Say we have a pool named cvpost at some remote site and we want to flock jobs to it from our pool named rmpost. What kind of networking is necessary? Do the execute hosts need a routable IP (NAT or real) for download and/or upload? What about the submit host and central manager?

- Job itself: submit host sends job to remote central manager?
- transfer input files: The submit host sends files to the execute host?
- transfer output files: The execute host sends files to the submit host?

ANSWER: These paths need to be open

- From local schedd to remote collectord on condor port 9618
- From remote negotiator and execute hosts to local schedd. Here the execute hosts can be NATed.
- From local shadow to remote starterd. Use CCB. It allows execute hosts to live behind firewall and be NATed.

### /tmp

```
executable = /bin/bash
arguments = "-c '/bin/date > /tmp/date'"
should_transfer_files = yes
transfer_output_files = /tmp/date
#transfer_output_files = tmp/date
queue
```

If I write to /tmp/date and set **transfer\_output\_files = /tmp/date** I get errors like

```
Error from slot1_4@nmpost040.aoc.nrao.edu: STARTER at 10.64.10.140 failed to send
file(s) to <10.64.10.100:9618>: error reading from /tmp/date: (errno 2) No such
file or directory; SHADOW failed to receive file(s) from <10.64.10.140:35386>
```

It works if I set **transfer\_output\_files = tmp/date**

### /dev/shm

```
executable = /bin/bash
arguments = "-c '/bin/date > /tmp/date'"
should_transfer_files = yes
transfer_output_files = /dev/shm/date
#transfer_output_files = dev/shm/date
queue
```

If I write to /dev/shm/date I get errors setting **transfer\_output\_files = /dev/shm/date**

```
Error from slot1_4@nmpost040.aoc.nrao.edu: STARTER at 10.64.10.140 failed to send
file(s) to <10.64.10.100:9618>: error reading from
/dev/shm/date: (errno 2) No
such file or directory; SHADOW failed to receive file(s) from
<10.64.10.140:41516>
```

If I write to /dev/shm/date I get errors setting **transfer\_output\_files = dev/shm/date**

```
Error from slot1_4@nmpost040.aoc.nrao.edu: STARTER at 10.64.10.140 failed to send
file(s) to <10.64.10.100:9618>: error reading from
/dev/shm/date: (errno 2)
No such file or directory; SHADOW failed to receive file(s) from
<10.64.10.140:40380>
```

ANSWER: these are known issues and not surprising. It's debatable weather they are bugs or not. The issue is the job is "done" by the time transfer\_output\_files is used and since the job is done the bindmounts for /tmp and /dev/shm(which is a little different) are gone.

### pro-active glideins

Need to investigate gliding in based on lack of free slots rather than idle jobs. Can one query HTCondor for a CARTA-shaped slot (core, mem, disk)?

ANSWER: Greg thinks this is a good idea and might be useful as a condor-week talk.

### condor\_off vs condor\_drain

a -peaceful option to condor\_drain might be perfect. Low priority for NRAO.

ANSWER: Yes condor\_drain is being worked on and this is one of the things.

### Transfer Plugins

Don't have condor block on the transfer plugin uploading. Low priority for NRAO.

ANSWER: This requires some serious work. Greb will ask Todd about it.

### More plugin woes

So let's say you have a plugin to transfer output files and this plugin fails because a destination directory, like nosuchdir, doesn't exist. All the plugin can do is indicate success or failure so it indicates failure. But that seems to cause HTCondor to disconnect/reconnect four times, the fail, then set the job to idle so it can try again later, which then disconnects/reconnects four times and ... Is there anything else the plugin can do to tell HTCondor to hold the job instead of restart?

```
executable = /bin/sleep
arguments = "127"
output = nosuchdir/condor_out.log
error = nosuchdir/condor_err.log
log = condor.log
should_transfer_files = YES
transfer_output_files = _condor_stdout
```

```
# output_destination = nraosync://$ENV(PWD)
+WantIOProxy = True
queue

If you set either output or error to a directory that doesn't exist like output = nosuchdir/condor_out.log, then when the job ends, HTCondor will put the job on hold with a message like the following in the condor.log

040 (5062.000.000) 2022-09-30 08:28:58 Finished transferring output files
...
007 (5062.000.000) 2022-09-30 08:28:58 Shadow exception!
    Error from slot1_2@mpost040.aoc.nrao.edu: STARTER at 10.64.10.140 failed to send file(s) to <10.64.10.100:9618>; SHADOW at 10.64.10.100 failed to write to file /users/krowe/htcondor/nraosync/dir/stdout.5062.log: (errno 2) No such file or directory
    13 - Run Bytes Sent By Job
    354 - Run Bytes Received By Job
...
012 (5062.000.000) 2022-09-30 08:28:58 Job was held.
    Error from slot1_2@mpost040.aoc.nrao.edu: STARTER at 10.64.10.140 failed to send file(s) to <10.64.10.100:9618>; SHADOW at 10.64.10.100 failed to write to file /users/krowe/htcondor/nraosync/dir/stdout.5062.log: (errno 2) No such file or directory
    Code 12 Subcode 2

But if you have set output_destination to use the nraosync plugin like so output_destination = nraosync://$ENV(PWD) then you get four disconnect/reconnect events followed by a shadow exception (see below). Then HTCondor sets the job to idle so it can try again instead of putting it on hold. I assume this is because it doesn't know why the job failed because there isn't really a mechanism for the plugin to tell it why.

040 (5061.000.000) 2022-09-30 08:23:20 Finished transferring output files
...
022 (5061.000.000) 2022-09-30 08:23:20 Job disconnected, attempting to reconnect
    Socket between submit and execute hosts closed unexpectedly
    Trying to reconnect to slot1_3@mpost040.aoc.nrao.edu <10.64.10.140:9618?addrs=10.64.10.140-9618&alias=mmpost040.aoc.nrao.edu&noUDP&sock=startrd_5631_f9e2>
...
023 (5061.000.000) 2022-09-30 08:23:20 Job reconnected to slot1_3@mmpost040.aoc.nrao.edu
    startrd address: <10.64.10.140:9618?addrs=10.64.10.140-9618&alias=mmpost040.aoc.nrao.edu&noUDP&sock=startrd_5631_f9e2>
    starter address: <10.64.10.140:9618?addrs=10.64.10.140-9618&alias=mmpost040.aoc.nrao.edu&noUDP&sock=slot1_3_5795_da05_282>
...
040 (5061.000.000) 2022-09-30 08:23:20 Started transferring output files
...
040 (5061.000.000) 2022-09-30 08:23:20 Finished transferring output files
...
022 (5061.000.000) 2022-09-30 08:23:20 Job disconnected, attempting to reconnect
    Socket between submit and execute hosts closed unexpectedly
    Trying to reconnect to slot1_3@mpost040.aoc.nrao.edu <10.64.10.140:9618?addrs=10.64.10.140-9618&alias=mmpost040.aoc.nrao.edu&noUDP&sock=startrd_5631_f9e2>
...
023 (5061.000.000) 2022-09-30 08:23:20 Job reconnected to slot1_3@mmpost040.aoc.nrao.edu
    startrd address: <10.64.10.140:9618?addrs=10.64.10.140-9618&alias=mmpost040.aoc.nrao.edu&noUDP&sock=startrd_5631_f9e2>
    starter address: <10.64.10.140:9618?addrs=10.64.10.140-9618&alias=mmpost040.aoc.nrao.edu&noUDP&sock=slot1_3_5795_da05_282>
...
040 (5061.000.000) 2022-09-30 08:23:20 Started transferring output files
...
040 (5061.000.000) 2022-09-30 08:23:21 Finished transferring output files
...
022 (5061.000.000) 2022-09-30 08:23:21 Job disconnected, attempting to reconnect
    Socket between submit and execute hosts closed unexpectedly
    Trying to reconnect to slot1_3@mpost040.aoc.nrao.edu <10.64.10.140:9618?addrs=10.64.10.140-9618&alias=mmpost040.aoc.nrao.edu&noUDP&sock=startrd_5631_f9e2>
...
023 (5061.000.000) 2022-09-30 08:23:21 Job reconnected to slot1_3@mmpost040.aoc.nrao.edu
    startrd address: <10.64.10.140:9618?addrs=10.64.10.140-9618&alias=mmpost040.aoc.nrao.edu&noUDP&sock=startrd_5631_f9e2>
    starter address: <10.64.10.140:9618?addrs=10.64.10.140-9618&alias=mmpost040.aoc.nrao.edu&noUDP&sock=slot1_3_5795_da05_282>
...
040 (5061.000.000) 2022-09-30 08:23:21 Started transferring output files
...
040 (5061.000.000) 2022-09-30 08:23:21 Finished transferring output files
...
022 (5061.000.000) 2022-09-30 08:23:21 Job disconnected, attempting to reconnect
    Socket between submit and execute hosts closed unexpectedly
    Trying to reconnect to slot1_3@mpost040.aoc.nrao.edu <10.64.10.140:9618?addrs=10.64.10.140-9618&alias=mmpost040.aoc.nrao.edu&noUDP&sock=startrd_5631_f9e2>
...
023 (5061.000.000) 2022-09-30 08:23:21 Job reconnected to slot1_3@mmpost040.aoc.nrao.edu
    startrd address: <10.64.10.140:9618?addrs=10.64.10.140-9618&alias=mmpost040.aoc.nrao.edu&noUDP&sock=startrd_5631_f9e2>
    starter address: <10.64.10.140:9618?addrs=10.64.10.140-9618&alias=mmpost040.aoc.nrao.edu&noUDP&sock=slot1_3_5795_da05_282>
...
040 (5061.000.000) 2022-09-30 08:23:21 Started transferring output files
...
040 (5061.000.000) 2022-09-30 08:23:21 Finished transferring output files
...
007 (5061.000.000) 2022-09-30 08:23:22 Shadow exception!
    Error from slot1_3@mmpost040.aoc.nrao.edu: Repeated attempts to transfer output failed for unknown reasons
    0 - Run Bytes Sent By Job
    354 - Run Bytes Received By Job

ANSWER: This is a bug. CHTC would like to implement better error handling here.

Workaround could be to set the following in the config file on the submit host. But this may be problematic on SSA's container submit host. It should cause condor to fail the job if the nosuchdir doesn't exist. I think its just best to note this as a bug and wait for CHTC to implement better error handling that allows the plugin to tell HTCondor how to fail. Or something like that.

• SUBMIT_SKIP_FILECHECKS=false
```

## DAG hosts

Is there a way to guarantee all the nodes of a DAG run on the same hostname without specifying the specific hostname?  
 An example would be the first node copies some data to local host storage, then all the other nodes read that data.  
 ANSWER: have the DAG post script figure out what hostname the node ran on and then modify or create the submit file for the next node.

## glidein memory requirements

Twice now I have doubled the memory for the pilot job. On Jul. 14, 2022 from 1GB to 2GB and just now (Aug. 25, 2022) from 2GB to 4GB. This is because the condor daemon like condor\_starter exceeded the memory and was OOM killed. This second time was during nraosync uploading files. Is there a suggested amount of memory for Slurm for glidein jobs?  
 ANSWER: The stand assumes it has control of all physical memory and don't check if they are in a group or not. If I run into this again, try and track down what is actually happening. Greg would like to know. He is surprised because the HTCondor daemons should only need MBs not GBs.

## no match found

When a job stays idle for a long time and its LastReMatchReason = "no match found" what are some good places to look to see why it isn't finding a match? For example, if you make a type-o and set the following (note the misspelling of nraosync)  
 transfer\_input\_files = \$ENV(HOME)/ssh/condor\_transfer, nraosync://\$ENV(PWD)/testdir  
 ANSWER: condor\_q -better  
 It doesn't know about quotas or fairshare or per-machine resource limits like +IOHeavy or other such adds.

## Accounts

Can Felipe get an account? Also, you might want to ask James if he still needs his account now that he no longer works for the NRAO.  
 ANSWER: done.

## Start glidein node if there isn't one free

Previously, our factory.sh script would start a glidein job in Slurm if there was a job idle. But now that we want jobs to start as quickly as possible, our factory.sh script now starts a glidein job if there isn't enough free resources available. To do this we had to define what "free resources" meant so we went with MIN\_SLOTS:=8 and MIN\_MEMORY:=16384 since we use dynamic slots. We also had to set a "default" machine add on all the nodes that we wanted available to this glidein job. This is so that the factory.sh script doesn't check nodes that are in the VCLASS or CVPOST or other such groups. I could have explicitly excluded those groups but that wouldn't scale well if we ever created more groups.  
 ANSWER: Greg thinks this is a perfectly crumulent way to do this.

## OS

What RHEL8-like OS is CHTC going to use or is using? CentOS8/Stream, Alma, Rocky, etc? Looks like CentOS8 Stream. Any thoughts?  
 ANSWER: yes they are using CentOS8/Stream. So far so good.

## PATH getting data to execute hosts

What are the preferred methods? http? nraosync? other?  
 ANSWER: http, s3 or other plugins. Or OSDF.  
<https://osg-htc.org/services/osdf.html>  
 OSDF (Open Science Data Federation) There are "data origins" which are basically webservers with cache. Long term we might be able to have our own data origin that authenticates to NRAO and shares data from our Lustre filesystem to their Ceph system via some way. This is like an object store so you can't really update files but you can make new ones.  
 /mnt/stash/ospool/PROTECTED/ Copy data from NRAO to this path and then you can access it in the job via  
 transfer\_input\_files = stash://ospool/PROTECTED/user/file  
 transfer\_output\_files = stash://ospool/PROTECTED/user/file  
[https://portal.osg-htc.org/documentation/htc\\_workloads/manager/osp\\_data/stashcache/](https://portal.osg-htc.org/documentation/htc_workloads/manager/osp_data/stashcache/)  
 I think this is cooler than using the nraosync plugin.

## PATH GPUs

We only see four GPUs on PATH right now. What is the timeline to get more? Does PATH flock to other sites with GPUs?  
 ANSWER: hosts may be dynamic and only come on-line as needed with some k8s magic. Christina is checking. Greg is pretty sure there should be way more than just 4 GPUs in PATH. PATH is made up of six different sites <https://path-cs.io/facility/index.html> each of these sites provides hardware.

## Disk Space

Since neither HTCondor nor cgroups control scratch space, how can we keep jobs from using up all the scratch space on a node and causing other jobs to fail as well?

ANSWER: specify a periodic hold in the standr. every 6 seconds standr can check and put the job on hold. Greg can look up the syntax. Someday, condor starter will make an efemeral filesystem (loopback) on the scratch area with the requested disk space. This is coming soon.  
<https://htcondor-wiki.cs.wisc.edu/index.cgi/wiki?title=HowToLimitDiskUsageOfJobs>

## Squid on PATH?

Does PATH use a squid server so that we can wget something once and have it cached for a while?

ANSWER: Greg thinks PATH does this as well.

## RADIAL workload balance

If there are only two users submitting jobs to a cluster, will HTCondor try to balance the workload between the two users? For example will it prioritize user2 jobs if user1 jobs are using the majority of resources? I think I read this about HTCondor's fair-share algorithm but I am not sure.

ANSWER: yes condor does this. There are knobs to adjust user priorities (user1 is twice the priority of user2, etc). You can also specify the length of the half life. There are many other ways to do something like this.

## Cluster domain names

This is not an HTCondor question but perhaps Greg has some insight. Let's say I am setting up a turn-key cluster. We deliver a rack of compute nodes with one head node. That one head node will need Internet access for SSH, DNS, etc. But the compute nodes don't need any Internet access. You can only get to them via ssh by first sshing to the head node and the only host the use as a name server is the head node. So, my question is what TLD should all these compute nodes be in? I would like to use a local, non-routable TLD analogous to non-routable IP ranges like 10.0.0.0/8 or 192.168.0.0/16. But there doesn't seem to be such a thing defined by ICANN.

- example is only for documentation.
- test is only for testing and our product is in production.
- invalid is meant for initial scripts that must be changed.
- *home.arpa* is for home networking (thank you IETF). Not ideal but better than the above.
- local is all tied up with mDNS thanks to Apple. But Kubernetes uses .local without mDNS.
- internal is used by both Google and Amazon which is pretty compelling.
- There have been many Internet-Draft documents since 2017 attempting to address this but none have become RFCs.

Have you heard of any clusters using any "private" TLDs? Can we just use IPs and not use names?

ANSWER: Greg looked at their nodes and saw both .local and .internal in use

## 2FA

I don't seem to be prompted for two-factor authentication when I login to CHTC. Should I?

ANSWER: Greg will ask. I'm not going to worry about it unless I can no longer login.

## Using containers at PATH

Anything special to do? Singularity/Apptainer?

ANSWER: Greg doesn't think so but Singularity would be the first to test.

## Docker at PATH

Doesn't seem to work. My docker universe job just stayed idle for three days.

ANSWER: since PATH uses docker containers in kubernetes, they probably can't support docker containers in a docker container. So Singularity is the answer.

## PATH libraries

It seems that not all the PATH execution hosts have the same set of libraries. For example my program, <https://github.com/axboe/fo.git> requires `libmvec.so.1`. It fails on all Expanse, SYRA GPU, UNL GPU nodes. It works on all WISC, FIU, SYRA non-GPU nodes.

Work on

- FIU-PATH-EP-osgvo-docker-pilot-gpu-949b9d958-57vzr (GPU node)
- FIU-PATH-EP-osgvo-docker-pilot-5578c6c47-6m9gl
- FIU-PATH-EP-osgvo-docker-pilot-5578c6c47-h2pwk
- FIU-PATH-EP-osgvo-docker-pilot-5578c6c47-hcbcp
- SYRA-PATH-EP-osgvo-docker-pilot-846a9d9556-8bump
- SYRA-PATH-EP-osgvo-docker-pilot-66f7756656-pm2Br
- UNL-PATH-EP-osgvo-docker-pilot-6b54769d-h622
- UNL-PATH-EP-osgvo-docker-pilot-579bd6994-p9jv
- WISC-PATH-EP-osgvo-docker-pilot-7b4d97d9f6-nqm6m

Fails on

- Expanse-PATH-EP-osgvo-docker-pilot-gpu-853b77b96-d2abd (GPU node)
- Expanse-PATH-EP-osgvo-docker-pilot-gpu-853b77b96-g2hnt (GPU node)
- Expanse-PATH-EP-osgvo-docker-pilot-76687d6f5c-5f4gx
- SYRA-PATH-EP-osgvo-docker-pilot-gpu-5bdc9969c-hfwht (GPU node)
- UNL-PATH-EP-osgvo-docker-pilot-gpu-847d757667-xnhwb (GPU node)

ANSWER: use singularity

## Flocking to PATH

Is flocking to PATH expected to be a thing either now or in the future?

Apparently there may be a way to setup PATH as an annex <https://htcondor.org/experimental/ospool/byoc/path-facility> if we knew our PROJECT\_ID

ANSWER: annex <https://htcondor.org/experimental/ospool/byoc/path-facility>

## gpu\_burn

When I run <https://github.com/wilicc/gpu-burn> on a PATH node, the nvidia-smi output only shows a couple of extra watts, about 400MiB of memory used and no GPU-Utilization. Normally, this program pegs GPUs at 100% utilization.

ANSWER: Greg agrees with krowe this might be because the GPUs are A100s

## Transfer Mechanism granularity?

Say you are transferring two files A and B. Is there a way to tell how long HTCondor took to transfer A and how long it took to transfer B? Are they even transferred serially? This may help Felipe tell how long it takes to transfer an MS vs the CFCache.

ANSWER: the files are transferred serially. And there is no way to tell file A vs file B.

## Excluding nodes

How can we tell HTCondor to not run on nodes like FIU\* or UNL\*?

requirements = GLIDEIN\_ResourceName != "FIU-PATH-EP"

ANSWER: this is probably the best solution. CHTC does have a taxonomy for naming things but it may not always get used.

## Job Shape Plots

Felipe has some plots showing what each subjob is doing in a large job (queue 32). Does HTCondor have or suggest any tools for plotting this sort of thing?

ANSWER: not really. Greg really likes Felipe's "gant charts". Matplotlib is in common use. Admins use Graphana. There are python modules for reading user logs.

## Nvidia MIG support?

Dividing a GPU into multiple (Max 8) GPU slices. This started with the Nvidia Ampere architecture. Does HTCondor support this?

ANSWER: Condor supports MIGs but no dynamically. The admin has to split the GPUs statically and then condor can use it.

## Upgrade to 10.x

What are the current recommendations for upgrading to 10.x? Order of upgrades? Mixed mode?

ANSWER: mixed mode should be fine.

## condor\_ssh\_to\_job

If I am vlapipe@nmpost-master and want to connect to a job submitted from vlapipe@hamilton (actually a container running on hamilton) it doesn't work. Should it? I am guessing it doesn't because in this case hamilton is a container and not the actual host hamilton.

vlapipe@nmpost-master\$ condor\_ssh\_to\_job name hamilton 3691

Failed to send GET\_JOB\_CONNECT\_INFO to schedd

ANSWER: This is less of an issue as Amy can use the wf\_inspector tool from SSA to connect to jobs.

## Preference

We want some jobs to run on a set of nodes (e.g. NMT VLASS), but if those aren't available, then run on the default set of nodes (e.g. DSOC VLASS).

I should be able to use the rank expression to do this, right? E.g.

Rank = (machine == "nmpost039.aoc.nrao.edu")

But when I run a job with this, it runs on nmpost038 instead of nmpost039. There is nothing wrong with nmpost039 and both nodes have the same config file. I can require the job to run on nmpost039 with the following

requirements = (machine == "nmpost039.aoc.nrao.edu")

HTCondor should select the most restrictive first, right? Is that NMT VLASS or DSOC VLASS.

Rank seems work at PATH but not NRAO nor CHTC.

ANSWER: the negotiator finds all slots that match. It sorts them by NEGOTIATOR\_PRE\_JOB\_RANK, then by JOB\_RANK, then by NEGOTIATOR\_POST\_JOB\_RANK. It works at PATH because the available nodes at PATH are all empty so NEGOTIATOR\_PRE\_JOB\_RANK is the same value for all of them so then it goes to the JOB\_RANK. We could set NEGOTIATOR\_POST\_JOB\_RANK == JOB\_RANK in the config file on the Central Manager, but that would not pack jobs the way we like.

By default NEGOTIATOR\_PRE\_JOB\_RANK = (1000000 \* MyRank) + (100000 \* (RemoteOwner ==? UNDEFINED)) - (100000 \* Cpus) - Memory

Since we disable preemption, I think MyRank, which is the rank in the machine classad, is always 0.0. What if we replaced it with Target.Rank which would be the rank in the job classad? By default rank in the job classad is also 0.0 unless the user sets rank in the submit description file.

So change it to NEGOTIATOR\_PRE\_JOB\_RANK = (10000000 \* Target.Rank) + (1000000 \* (RemoteOwner ==? UNDEFINED)) - (100000 \* Cpus) - Memory in the config file on the Central Manager.

## condor\_q output totals

The **Total** lines at the end of **condor\_q** are the same format at NRAO and PATH. Why and how?

```
Total for query: 0 jobs; 0 completed, 0 removed, 0 idle, 0 running, 0 held, 0 suspended
Total for krowe: 0 jobs; 0 completed, 0 removed, 0 idle, 0 running, 0 held, 0 suspended
Total for all users: 0 jobs; 0 completed, 0 removed, 0 idle, 0 running, 0 held, 0 suspended
```

But at CHTC it is **only one line**.

```
0 jobs; 0 completed, 0 removed, 0 idle, 0 running, 0 held, 0 suspended
```

ANSWER: Don't know. But this is pretty minor and not really a problem just a curiosity.

## ypbind.service in condor.service

Pretty minor but I noticed that the systemd unit file for condor indicates it needs to start after ypbind. I can see this on CHTC. Even NRAO is getting rid of ypbind. Y'all can probably remove this.

```
submit2 nu_kscott >systemctl cat condor | grep After
After=network-online.target nslcd.service ypbind.service time-sync.target nfs-client.target autofs.service
```

K. Scott will look into a proper solution. perhaps systemd has a directory service definition to use.

ANSWER: ypbind has largely been replaced by sssd which can use various directory services (LDAP, Active Directory, IdMIPA, etc)

## Time Zones

We have multiple submit hosts and some of them are in different time zones. Is there a condor command that could tell me the timezone of a particular submit host? Something like **condor\_status -i -schedd hamilton | grep -i zone** Logging into some of these submit hosts is hard because they are containers. I need to know the timezone so that I can cross reference the logs correctly.

ANSWER: No. you can advertise whatever value you want in the classad. We could create a timezone variable. Greg also thinks there may be a `__hacky__` way to do this.

## Time zone logging?

Is there a way to know what time zone a set of logs are in? Or can you change the time zone that your condor job uses so that they are always in UTC?

ANSWER: Not at the user level. You can set **DEFAULT\_USERLOG\_FORMAT\_OPTIONS = UTC** in the condor config file (not submit file).

## PATH credits

Is there a way to check your credit status at PATH?

ANSWER: Greg doesn't know. Ask Christina.

## ARM Mac support (dlyons)

A problem we're about to have is, how do we submit jobs from an ARM Mac container?

My sense is that the only way to do that is to run the container image as if it were a submit host. And that there are going to be two problems with this:

1. There is no ARM build of any of the Condor docker images
2. There may be architecture mismatch issues

I am willing to try to build Condor images for them but I have found their documentation for building containers somewhat obtuse. If they have a cheat sheet I can steal from that would be useful.

Anyway, I don't really know if there is another approach, but if there is, that might be better.

If you want me at the meeting to ask about this, I can come.

ANSWER:

<https://www-auth.cs.wisc.edu/lists/hiccondor-users/2023-March/mso00130.shtml>

CHTC uses QEMU in docker to build their non-x86\_64 binaries. It's slow but surprisingly reliable.

## Flocking to radial

I can flock from testpost cluster to nmpost cluster by just setting **FLOCK\_TO = nmpost-cm-vml.aoc.nrao.edu** and copying the idtoken from nmpost. But when I try flocking from testpost cluster to radial cluster instead of nmpost using the same procedure, it doesn't work. I see the jobs in condor\_q on radialhead but the jobs just stay idle. I don't see any errors in any logs or any indication why it isn't working.

Testpost and nmpost are both HTCondor-9 with the same user space, FILESYSTEM\_DOMAIN, and UID\_DOMAIN. Also, the radial cluster is on an isolated network (192.168.0.0/24) and the central manager is multi-homed.

What is the data path when a job flocks? Does the flocked\_to startd have to contact the original schedd or the flocked\_to schedd?

ANSWER: The schedd advertises the demand to the remote pool. Greg will need to think about this.

2023-03-29 krowe: I upgraded testpost to HTCondor-10 and tried flocking to radial again. Now I see the following in the SchedLog on testpost-master. So it looks like the original schedd needs to contact the eventual startd. That's going to be a problem because the startd is on an isolated network.

```
03/29/23 08:30:28 (pid:119154) attempt to connect to <192.168.0.32:9618> failed: timed out after 45 seconds.
03/29/23 08:30:28 (pid:119154) Failed to send REQUEST_CLAIM to startd slot1@radial001.nrao.radial.local <192.168.0.32:9618?addr=192.168.0.32:9618&alias=radial001.nrao.radial.local&noUDP&sock=startd_1460_045c> for krowe: SECMAN:2003:TCP connection to startd slot1@radial001.nrao.radial.local <192.168.0.32:9618?
addr=192.168.0.32:9618&alias=radial001.nrao.radial.local&noUDP&sock=startd_1460_045c> for krowe failed.
03/29/23 08:30:28 (pid:119154) Match record (slot1@radial001.nrao.radial.local <192.168.0.32:9618?addr=192.168.0.32:9618&alias=radial001.nrao.radial.local&noUDP&sock=startd_1460_045c> for krowe, 773.0) deleted
```

To test this I made radialhead (which is dual-homed) a startd and I can flock from testpost-master to radialhead. So how do I run a job remotely on a startd on an isolated network?

- I don't think CCB will work because the startd isn't NATed
- Annex?
- condor\_remote\_cluster?
- condor\_submit\_remote?

ANSWER: HTCondor really wants the execute host to at least have outbound networking.

ANSWER: HTCondor-c is source level routing from one schedd to another schedd. schedd-a is in NM and schedd-b is radial. With condor-c say submit this job but don't run it local run it on schedd-b. If you have multiple sites you have to select which site to use when you set the job.

## Object Store

What do people use to put files in stash? Tar? HDF5? zip? other?

ANSWER: HTCondor sets OMP\_NUM\_THREADS=1 which may affect the speed of uncompressing. Should test.

ANSWER: gzip can handle directories? But Greg thinks it may not matter. Tar is probably fine

QUESTION: what is system vs user time? I/O time?

KROWE: I unset OMP\_NUM\_THREADS and did manual tests outside of HTCondor using the NVMe on testpost001.

Uncompressed tar file

```
./usr/bin/time -f "real %e\nuser %U\nkernel %S\nwaits %w" tar xf hudef_n1.ms.tar
real 6.53
user 0.52
kernel 5.98
waits 85
```

Compressed tar file with gzip

```
testpost001 krowe >./usr/bin/time -f "real %e\nuser %U\nkernel %S\nwaits %w" tar xf gzip-6.tgz
real 36.97
user 35.58
kernel 6.65
waits 172421
```

Compressed tar file with bzip2

```
testpost001 krowe >./usr/bin/time -f "real %e\nuser %U\nkernel %S\nwaits %w" tar xf bzip2.tgz
real 311.63
user 310.06
kernel 16.43
```

| waits 1128218

ANSWER: Greg is as surprised as we are.

## Federation

Does this look correct?

<https://staff.nrao.edu/wiki/bin/view/NM/HTCondor-federations>

ANSWER: yes

## Reservations

Reservations from the Double Tree were for Sunday Jul.9 through Thursday Jul. 13 (4 nights). But I need at least until Friday Jul. 14 right?

ANSWER: Greg will look into it.

## Scatter/Gather problem

At some point we will have a scatter/gather problem. For example we will launch 10\*5 jobs, each of which will produce some output that will need to be summed with the output of all the other jobs. Lanching 10\*5 jobs is not hard. Dong the sumation is not hard. Moving all the output around is the hard part.

One idea is to have a dedicated process running to which each job uploads its output. This process could sum output as it arrives; it doesn't need to wait until all the output is done. It would be nice if this process also ran in the same HTCondor environment (PATH, CHTC, etc) because that would keep all the data "close" and presumably keep transfer times short.

ANSWER: DAGs and sub-DAGs of course. Provisioner node was created to submit jobs into the cloud. It exists as long as the DAG is working.

Nodes talking to each other becomes difficult in federated clusters.

<https://cdl.cse.nrl.edu/software/taskvine/>

makeflow

Astra suggests something like Apache Beam for ngVLA data which is more of a data approach than a compute approach.

## What is annex?

Yet another way to federate condor clusters. Annexes are useful when you have an allocation on a system (AWS) you have the ability to start jobs on a system. You give annex your credentials and how many workers you want. Annex will launch the startds and create a local central manager. It then configures flocking from your local to the remote pool. So in a sense annex is an ephemeral flocking relationship for just the one person setting up the annex.

## Condor and Kubernetes (k8s)

Condor supports Docker, Singularity, and Apptainer. In OSG the Central Managers are in K8s and most of the Schedds are also in k8s. In the OSPool some worker nodes are in k8s and they are allowed to run unprivileged Apptainer but not Docker (because privileges). PATH has worker nodes in k8s. They are backfilled on demand.

## NSF Archive Storage

Are you aware of any archive storage funded by NSF we could use? We are looking for off-site backup of our archive (NGAS).

ANSWER: Greg doesn't know of one.

## Hung Jobs and viewing stdout

We have some jobs that seem to hang possibly because of a race condition or whatnot. I'm pretty sure it is our fault. But, the only way I know to tell is to login to the node and look at \_condor\_stdout in the scratch area. That gets pretty tedious when I want to check hundreds of jobs to see which ones are hung. Does condor have a way to check the \_condor\_stdout of a job from the submit host so I can do this programatically?

I thought **condor\_tail** would be the solution but it doesn't display anything.

ANSWER: condor\_ssh\_to\_job might be able to be used non-interactively. I will try that.

ANSWER: use the FULL jobid with condor\_tail. E.g. **condor\_tail 12345.0** Greg has submitted a patch so you don't have to specify the Procid (J.O).

## Bug: condor\_off -peaceful

```
testpost-cm-vml root >condor_off -peaceful -name testpost002
Sent "Set-Peaceful-Shutdown" command to startd testpost002.aoc.nrao.edu
Can't find address for schedd testpost002.aoc.nrao.edu
Can't find address for testpost002.aoc.nrao.edu
Perhaps you need to query another pool.
```

Yet it works without the -peaceful option

```
testpost-cm-vml root >condor_off -name testpost002
Sent "Kill-All-Daemons" command to master testpost002.aoc.nrao.edu
```

ANSWER: Add the **-startd** option. E.g. condor\_off -peaceful -startd -name <hostname> Greg thinks it might be a regression (another bug). This still happens even after I set all the CONDOR\_HOST knobs to [testpost-cm-vml.aoc.nrao.edu](#). So it is still a bug and not because of some silly config I had at NRAO.

## File Transfer Plugins and HTCondor-C

Is there a way I can use our nraosync plugin on radial001? Or something similar?

SOLUTION: ssh tunnels

## Condor Week (aka Throughput Week)

July 10-14, 2023. Being co-run with the OSG all hands meeting. At the moment, it is not hybrid but entirely in-person. <https://path-cc.io/htc23>

## PROVISIONER node

When I define a PROVISIONER node, that is the only node that runs. The others never run. Also, the PROVISIONER job always returns 1 "exited normally with status 1" even though it is just running /bin/sleep.

```
JOB node01 node01.htc
JOB node02 node02.htc
JOB node03 node03.htc
PARENT node01 CHILD node02 node03
PROVISIONER prov01 provisioner.htc
```

ANSWER: my prov01 job needs to indicate when it is ready with something like the following but that means the provisioner job has to run in either the local or scheduler universes because our execute nodes cant run condor\_qedit.

```
condor_qedit myJobId ProvisionerState 2
```

But execute hosts cant run condor\_qedit so this really only works if you set universe to local or scheduler.

## Does CHTC have resources available for VLASS?

Our Single Epoch jobs

- Are parallelizable with OpenMPI
- 64GB of memory
- ~150GB of storage
- can take week(s) to run
- no checkpointing
- Looking at using GPUs but not there yet

Brian was not scared by this and gave us a form to fill out

<https://chtc.cs.wisc.edu/uw-research-computing/form.html>

ANSWER: Yes. We and Mark Lacy have started the process with CHTC for VLASS.

## Annex to PATH

<https://htcondor.org/experimental/ospool/byoc/path-facility>

ANSWER: Greg doesn't know but he can connect me with someone who does.

Tod Miller is the person to ask about this.

## Hold jobs that exceed disk request

ANSWER: <https://htcondor-wiki.cs.wisc.edu/index.cgi/wiki?p=HowToLimitDiskUsageOfJobs>

### condor\_userprio

We want a user (vlapipe) to always have higher priority than other users. I see we can set this with condor\_userprio but is that change permenent?  
ANSWER: There is no config file for this. Set the priority\_factor of vlapipe to 1. That is saved on disk and should persist through reboots and upgrades.

### Submitting jobs as other users

At some point in the future we will probably want the ability for a web process to launch condor jobs as different users. The web process will probably not be running as root. Does condor have a method for this or should we make our own setuid root thingy? Tell dlyons the answer.  
ANSWER: HTCondor doesn't have anything for this. So it is up to us to do some sud-fu.

### SSH keys with Duo

I tried following the link below to setup ssh such that I don't have to enter both my password and Duo every time I login to CHTC. It doesn't create anything in ~/.ssh/connections after I login. Thoughts?  
<https://chtc.cs.wisc.edu/uw-research-computing/configure-ssh>  
ANSWER: Greg doesn't know what to do here. We should ask Christina.

### HTCondor-C and requirements

submitting jobs from the container on shipman as vlapipe to the NRAO RADIAL prototype cluster seems to ignore requirements like the following. Is this expected?

```
| requirements = (machine == "radial001.nrao.radial.local")
```

and

```
| requirements = (VLASS == True)
| +partition = "VLASS"
```

It also seems to ignore things like

```
| request_memory = 100G
```

ANSWER:

[https://chtc.concord.readthedocs.io/en/latest/cond-computing/cond-universe.html#highlight=remote\\_requirements%5Bhtcondor-c-job-submission](https://chtc.concord.readthedocs.io/en/latest/cond-computing/cond-universe.html#highlight=remote_requirements%5Bhtcondor-c-job-submission)

But I am still having problems.

This forces job to run on radial001

```
| +remote_Requirements = ((machine == "radial001.nrao.radial.local"))
```

This runs on radialhead even though it only has 64G

```
| +remote_RequestMemory = 102400
```

This runs on radialhead even though it doesn't have a GPU

```
| request_gpus = 1
| +remote_RequestGPUs = 1
```

ANSWER: This works

```
| +remote_Requirements = ((machine == "radial001.nrao.radial.local") && memory > 102400)
```

as does this

```
| +remote_Requirements = (memory > 102400)
```

but Greg will look into why +remote\_RequestMemory doesn't work. It should.

### Select files to transfer dynamically according to job-slot match

We currently have separate builds of our GPU software for CUDA Capability 7.0 and 8.0, and our jobs specify that both builds should be transferred to the EP, so that the job executable selects the appropriate build to run based on the CUDA Capability of the assigned GPU. Is there a way to do this selection when the job is matched to a slot, so that only the necessary build is transferred according to the slot's CUDA Capability?

ANSWER: \$\$() means expand this locally from the jobad. \$\$() means expand at job start time.

```
executable = my_binary$$$(GPUs_capability)
```

```
executable = my_binary$$$(int(GPUs_capability)) # Felipe said this actually works
```

```
executable = my_binary$$$( classad_express(GPUS_capability) ) # Hopefully you don't need this
```

### CPU/GPU Balancing

We have 30 nodes in a rack at NMT with a power limit of 17 kW and we are able to hit that limit when all 720 cores (24 cores \* 30 nodes) are busy. We want to add two GPUs to each node but that would almost certainly put us way over the power limit if each node had 22 cores and 2 GPUs busy. So is there a way to tell HTCondor to reserve X cores for each GPU? That way we could balance the power load.

JOB TRANSFORMS work per schedd so that wouldn't work on the startd side which is what we want.

IDEA: **NUM\_CPUS = 4** or some other small number greater then the number of GPUs but limiting enough to keep the power draw low.

ANSWER: There isn't a knob for this in HTCondor but Greg is interested in this and will look into this.

WORKAROUND: MODIFY\_REQUEST\_EXPR\_REQUESTCPUS may help by setting each job gets 8cores or something like.

```
| MODIFY_REQUEST_EXPR_REQUESTCPUS = quantize(RequestCpus, isUndefined(RequestGpus) ? [1] : [8, 16, 24, 32, 40])
```

That is, when a job comes into the startd, if it doesn't request any GPUs, allocate exactly as many cpu cores as it requests. Otherwise, allocate 8 times as many cpus as it requests.

This seems to work. If I ask for 0 GPUs and 4 CPUs, I am given 0 GPUs and 4 CPUs. If I ask for 1 GPU and don't ask for CPUs, I am given 1 GPU and 8 CPUs.

But if I ask for 2 GPUs and don't ask for CPUs, I still am only given 8 CPUs. I was expecting to be given 16 CPUs. This is probably fine as we are not planning on more than 1 GPU per job.

But if I ask for 1 GPU and 4 CPUs, I am given 1 GPU and 8 CPUs. That is probably acceptable.

2024-01-24 krower: Assuming a node can draw up to 550 Watts when all 24 cores are busy and that node only draws 150 Watts when idle, and that we have 17,300 Watts available to us in an NMT rack.

- we should only need to reserve 3 cores per GPU in order to offset the 72 Watts of an Nvidia L4 GPU.
  - This would waste 60 cores.
- Or at least I suggest starting with that and seeing what happens. Another alternative is we just turn off three nodes if we put one L4 in each node.
  - This would waste 72 cores.

### Upgrading

CHTC just upgrades to the latest version when it becomes available, right? Do you ever run into problems because of this? We are still using version 9 because I can't seem to schedule a time with our VLASS group to test version 10. Much less version 23.

ANSWER: yes. The idea is that CHTC's users are helping them test the latest versions.

### Flocking to CHTC?

We may want to run VLASS jobs at CHTC. What is the best way to submit locally and run globally?

ANSWER: Greg thinks flocking is the best idea.

This will require 9618 open to nmpost-master and probably a static NAT and external DNS name.

### External users vs staff

We are thinking about making a DMZ ( I don't like that term ) for observers. Does CHTC staff use the same cluster resources that CHTC observers (customers) use?

ANSWER: There is no airgap at CHTC everyone uses the same cluster. Sometime users use a different AP but more for load balancing than security. Everyone does go through 2FA.

### Does PATH Cache thingy(tm) (a.k.a. Stash) work outside of PATH?

I see HTCondor-10.x comes with a stash plugin. Does this mean we could read/write to stash from NRAO using HTCondor-10.x?

ANSWER: Greg thinks you can use stash remotely, like at our installation of HTCondor.

### Curl\_ plugin doesn't do FTP

None of the following work. They either hang or produce errors. They work on the shell command line, except at CHTC where the squid server doesn't seem to grok FTP.

```
transfer_input_files = http://demo:password@test.rebex.net/readme.txt
transfer_input_files = http://ftp.gnu.org/welcome.msg
transfer_input_files = http://ftp.gnu.org/welcome.msg
transfer_input_files = http://ftp.slackware.com/welcome.msg
transfer_input_files = http://ftp.slackware.com/welcome.msg
```

2024-02-05: Greg thinks this should work and will look into it.

ANSWER: 2024-02-06 Greg wrote "Just found the problem with ftp file transfer plugin. I'm afraid there's no easy workaround, but I've pushed a fix that will go into the next stable release. "

### File Transfer Plugins and HTCCondor-C

I see that when a job starts, the execution point (radial001) uses our nrarsync plugin to download the files. This is fine and good. When the job is finished, the execution point (radial001) uses our nrarsync plugin to upload the files, also fine and good. But then the RADIAL schedd (radialhead) also runs our nrarsync plugin to upload files. This causes problems because radialhead doesn't have the \_CONDOR\_JOB\_AD environment variable and the plugin dies. Why is the remote schedd running the plugin and is there a way to prevent it from doing so?

Greg understands this and will ask the HTCCondor-c folks about it.

Greg thinks it is a bug and will talk to our HTCCondor-C peopole.

2023-08-07: Greg said the HTCCondor-C people agree this is a bug and will work on it.

2023-09-25 krowe: send Greg my exact procedure to reproduce this.

2023-10-02 krowe: Sent Greg an example that fails. Turns out it is intermittent.

2024-01-22 krowe: will send email to the condor list

ANSWER: It was K. Scott all along. I now have HTCCondor-C working from nrmpost and testpost clusters to the radial cluster using my nrarsync plugin to transfer both input and output files. The reason the remote AP (radialhead) was running the nrarsync plugin was because I defined it in the condor config like so.

```
FILETRANSFER_PLUGINS = $(FILETRANSFER_PLUGINS), /usr/libexec/condor/nrarsync_plugin.py
```

I probably did this early in my HTCCondor-C testing not knowing what I was doing. I commented this out, restarted condor, and now everything seems to be working properly.

### Quotes in DAG VARS

I was helping SSA with a syntax problem between HTCCondor-9 and HTCCondor-10 and I was wondering if you had any thoughts on it. They have a dag with lines like this

```
JOB SqDeg2/J232156-603000 split_condor
VARS SqDeg2/J232156-603000 jobname="$(JOB)" split_dir="SqDeg2/J232156+603000"
```

Then they set that split\_dir VAR to a variable in the submit description file like this

```
SPLIT_DIR = "${split_dir}"
```

The problem seems to be the quotes around \$(split\_dir). It works fine in HTCCondor-9 but with HTCCondor-10 they get an error like this in their pims\_split.dag.dagman.out file

```
02/28/24 16:26:02 submit error: Submit-1:Unexpected characters following doublequote. Did you forget to escape the double-quote by repeating it? Here is the quote and trailing characters: "SqDeg2/J232156+603000"
```

Looking at the documentation <https://htcondor.readthedocs.io/en/latest/version-history/its-versions-10-0.html#version-10-0-0> its clear they shouldn't be putting quotes around \$(split\_dir). So clearly something changed with version 10. Either a change to the syntax or, my guess, just a stricter parser.

Any thoughts on this?

ANSWER: Greg doesn't know why this changed but thinks we are now doing the right thing.

### OSDF Cache

Is there a way to prefer a job to run on a machine where the data is cached?

ANSWER: There is no knob in HTCCondor for this but CHTC would like to add one for this. they would like to glide in OSDF caches like they glide in nodes. But this is all long-term ideas.

### GPU names

HTCondor seems to have short names for GPUs which are the first part of the UUID. Is there a way to use/get the full UUID? This would make it consistent with nvidia-smi.

ANSWER: Greg thinks you can use the full UUID with HTCCondor.

But cuda\_visible\_devices only provides the short UUID name. Is there a way to get the long UUID name from cuda\_visible\_devices?

ANSWER: You can't use id 0 because 0 will always be the first GPU that HTCCondor chose for you. Some new release of HTCCondor supports NVIDIA\_VISIBLE\_DEVICES which should be the full UUID.

### Big Data

Are we alone in needing to copy in and out many GBs per job? Do other institutions have this problem as well? Does CHTC have any suggestions to help? Sanja will ask this of Bockleman as well.

ANSWER: Greg thinks our transfer times are not uncommon but our processing time is shorter than many. Other jobs have similar data sizes. Some other jobs have similar transfer times but process for many hours. Maybe we can constrain our jobs to only run on sites that seem to transfer quickly. Greg is also interested in why some sites seem slower than others. Is that actually site specific or is it time specific or...

Felipe does have a long list of excluded sites in his run just for this reason. Greg would like a more declarative solution like "please run on fast transfer hosts" especially if this is dynamic.

### GPUs\_Capability

We have a host (testpost001) with both a Tesla T4 (Capability=7.5) and a Tesla L4 (Capability=8.9) and when I run condor\_gpu\_discovery -prop I see something like the following

```
DetectedGPUs:"GPU-ddc998f9 GPU-40331b00"
Common=[ DriverVersion=12.20; ECCEnabled=true; MaxSupportedVersion=12020; ]
GPU_40331b00=[ id="GPU-40331b00"; Capability=7.5; DeviceName="Tesla T4"; DevicePciBusId="0000:3B:00.0"; DeviceUuid="40331b00-c3b6-fa9a-8f8d-33bec2fcd29c"; GlobalMemoryMb=14931; ]
GPU_ddc998f9=[ id="GPU-ddc998f9"; Capability=8.9; DeviceName="NVIDIA L4"; DevicePciBusId="0000:5E:00.0; DeviceUuid="ddc998f9-99e2-d9c1-04e3-7cc023a2aa5f"; GlobalMemoryMb=22491; ]
```

The problem is 'condor\_status -compact -constraint 'GPUs\_Capability >= 7.0' doesn't show testpost001. It does show testpost001 when I physically remove the T4.

Requesting a specific GPU with 'RequireGPUs = (Capability >= 8.0)' or 'RequireGPUs = (Capability <= 8.0)' does work however so maybe this is just a condor\_status issue.

We then replaced the L4 with a second T4 and then GPUs\_Capability functioned as expected.

Can condor handle two different capabilities on the same node?

ANSWER: Greg will look into it. They only recently added support for different GPUs on the same node. So this is going to take some time to get support in condor\_status. Yes this is just a condor\_status issue.

### Priority for Glidein Nodes

We have a factory.sh script that glides in Slurm nodes to HTCondor as needed. The problem is that HTCondor then seems to prefer these nodes to the regular HTCondor nodes such that after a while there are several free regular HTCondor nodes, and three glide-in nodes. Is there a way to set a lower priority on glide-in nodes so that HTCondor only chooses them if the regular HTCondor nodes are all busy? I am going to offline the glide-in nodes to see if that works but that is a manual solution not and automated one.

I would think NEGOTIATOR\_PRE\_JOB\_RANK would be the trick but we already set that on the CMs to the following so that RANK expressions in submit description files are honored and negotiation will prefer NMT nodes over DSOC nodes if possible.

```
NEGOTIATOR_PRE_JOB_RANK = (10000000 * Target.Rank) + (1000000 * (RemoteOwner !=?= UNDEFINED)) - (100000 * Cpus) - Memory
ANSWER: NEGOTIATOR_PRE_JOB_RANK = (10000000 * Target.Rank) + (1000000 * (RemoteOwner !=?= UNDEFINED)) - (100000 * Cpus) - Memory + 100000 * (site == "not-slurm")
```

I don't like setting not-slurm in the dedicated HTCondor nodes. I would rather set something like "glidein=true" or "glidein=1000" in the default 99-nrao config file and then remove it for the 99-nrao config in snapshots for dedicated HTCondor nodes. But that assumes that the base 99-nrao is for NM. Since we are sharing an image with CV we can't assume that. Therefore every node, whether dedicated HTCondor or not, will need a 99-nrao in its snapshot area.

SOLUTION

This seems to work. If I set **NRAOGLIDEIN = True** on a node, then that node will be chosen last. You may ask why not just add 100000000 ASTERISK (NRAOGLIDEIN == True). If I did that I would have to also set it to false on all the other nodes otherwise the negotiator would fail to parse NEGOTIATOR\_PRE\_JOB\_RANK into a float. So I check if it isn't undefined then check if it is true. This way you could set NRAOGLIDEIN to False if you wanted.

```
NEGOTIATOR_PRE_JOB_RANK = (10000000 * Target.Rank) + (1000000 * (RemoteOwner !=?= UNDEFINED)) - (100000 * Cpus) - Memory - 10000000 * ((NRAOGLIDEIN !=?= UNDEFINED) && (NRAOGLIDEIN == True))
```

I configured our pilot.sh script to add the NRAOGLIDEIN = True key/value pair to a node when it glides in to HTCondor. That is the simplest and best place to set this I think.

### K8s kubernetes

2024-04-15 krowe: There is a lot of talk around NRAO about k8s these days. Can you explain if/how HTCCondor works with k8s? I'm not suggesting we run HTCCondor on top of k8s but I would like to know the options.

Condor and k8s have different goals. Condor an infinite number of jobs for finite time each job. k8s runs a finite number of services for infinite time.

There is some support in k8s to run batch jobs but it isn't well formed yet. Running the condor services like the CM in k8s can make some sense.

The new hotness is using EBPf to change routing tables.

### RedHat8 Only

Say we have a few RedHat8 nodes and we only want jobs to run on those nodes that request RedHat8 with

```
requirements = (OpSysAndVer == "RedHat8")
```

I know I could set up a partition like we have done with VLASS but since HTCCondor already has an OS knob, can I use that?

Setting RedHat8 in the job requirements guarantees the job will run on a RedHat8 node, but how do I make that node not run jobs that don't specify the OS they want?

The following didn't do what I wanted.

```
START = $(START) && (TARGET.OpSysAndVer != "RedHat8")
```

Then I thought I needed to specify jobs where OpSysAndVer is not Undefined but that didn't work either. Either of the following do prevent jobs that don't specify an OS from running on the node but they also prevent jobs that DO specify an OS via either OpSysAndVer or OpSysMajorVer respectively.

```
START = $(START) && (TARGET.OpSysAndVer isnt UNDEFINED)
START = $(START) && (TARGET.OpSysMajorVer isnt UNDEFINED)
```

A better long-term solution is probably for our jobs (VLASS, VLA calibration, ingestion, etc) to ask for the OS that they want if they care. Then they can test new OSes when they want and we can upgrade OSes at our schedule (to a certain point). I think asking them to start requesting the OS they want now is not going to happen but maybe by the time RedHat8 is an option they and we will be ready for this.

ANSWER: unpars takes a classad expression and turns into a string then use a regex on it looking for opsysandver.

Is this the right syntax? Probably not as it doesn't work

```
START = $(START) && (regex("RedHat8.*", unpars(TARGET.Requirements)))
```

Greg thinks this should work. We will poke at it.

The following DOES WORK in the sense that it matches anything.

```
START = $(START) && (regex(".", unpars(TARGET.Requirements)))
```

None of these work

```
START = $(START) && (regex("RedHat8.*", unpars(Requirements)))
START = $(START) && (regex(".*", unpars(Requirements)))
START = $(START) && (regex("(OpSysAndVer.*", unpars(Requirements)))
START = $(START) && (regex("(OpSysAndVer.*", unpars(TARGET.Requirements)))
START = $(START) && (regex("(OpSysAndVer.*", unpars(Requirements)))
START = $(START) && (regex("(RedHat8.*", unpars(Requirements)))
START = $(START) && (regex("RedHat8", unpars(Requirements, "T"))
START = $(START) && (regex(".*RedHat8.*", unpars(Requirements, "T"))
START = $(START) && (regex(".*RedHat8.*", unpars(Requirements, "m"))
START = $(START) && (regex("OpSysAndVerIs.*", unpars(Requirements)))
START = $(START) && (regex("OpSysAndVerIs.*", unpars(Requirements)))
#START = $(START) && debug(regex("RedHat8.*", unpars(TARGET.Requirements)))
```

This should also work

in the config file

```
START = $(START) && target.WantToRunOnRedHat8Only
```

Submit file

```
MyWantToRunOnRedHat8Only = true
```

But I would rather not have to add yet more attributes to the EPs. I would like to use the existing OS attribute that HTCondor provides.

Wasn't there a change to PCRE to PCRE2 or something like that? Could that be causing the problem? 2023-11-13 Greg doesn't think so.

2024-01-03 krowe: Can we use a container like this? How does PATH do this?

```
+SingularityImage = "cvmfs/singularity.opensciencegrid.org/opensciencegrid/osgvo-el7.latest"
```

## See retired nodes

2024-04-15 krowe: Say I set a few nodes to offline with a command like **condor\_off-startd -peaceful -name nmpost120** How can I later check to see which nodes are offline?

- condor\_status -offline** returns nothing
- condor\_status -long nmpost120** returns nothing about being offline
- The following shows nodes where startd has actually stopped but it doesn't show nodes that are set offline but still running jobs (e.g. Retiring)
  - condor\_status -master -constraint 'STARTD\_StartTime == 0'**
- This shows nodes that are set offline but still running jobs (a.k.a. Retiring)
  - condor\_status |grep Retiring**

ANSWER: 2022-06-27

condor\_status -const Activity == "Retiring"

offline ads, which is a way for HTCondor to update the status of a node after the startd has exited.

condor\_drain -peaceful # CHTC is working on this. I think this might be the best solution.

Try this: **condor\_status -constraint 'PartitionableSlot && Cpus && DetectedCpus && State == "Retiring"**

or this: **condor\_status -const 'PartitionableSlot && State == "Retiring" -af Name DetectedCpus Cpus**

or: **condor\_status -const 'PartitionableSlot && Activity == "Retiring" -af Name Cpus DetectedCpus**

or: **condor\_status -const 'partitionableSlot && Activity == "Retiring" && cpus == DetectedCpus'**

None of which actually show nodes that have drained. I.e. were in state Retiring and are now done running jobs.

ANSWER: This seems to work fairly well. Not sure if it is perfect or not **condor\_status -master -constraint 'STARTD\_StartTime == 0'**

## Condor\_reboot?

Is there such a thing? Slurm has a nice one 'scontrol reboot HOSTNAME'. I know it might not be the condor way, but thought I would ask.

ANSWER: [https://htcondor.readthedocs.io/en/latest/admin-manual/configuration-macros.html#MASTER\\_SHUTDOWN\\_%3CName%3E](https://htcondor.readthedocs.io/en/latest/admin-manual/configuration-macros.html#MASTER_SHUTDOWN_%3CName%3E) and [https://htcondor.readthedocs.io/en/latest/man-pages/condor\\_set\\_shutdown.html](https://htcondor.readthedocs.io/en/latest/man-pages/condor_set_shutdown.html) maybe do the latter and then the former and possibly combined with condor\_off -peaceful. I'll need to play with it when I feel better.

## Felipe's code

Felipe to share his job visualization software with Greg and maybe present at Throughput 2024.

[https://github.com/ARDG-NRAO/LIBRA/tree/main/frameworks/htclean/read\\_htclean\\_logs](https://github.com/ARDG-NRAO/LIBRA/tree/main/frameworks/htclean/read_htclean_logs)

## Versions and falling behind

We are still using HTCondor-10.0.2. How far can/should we fall behind before catching up again?

ANSWER: Version 24 is coming out around condor week in 2024. It is suggested to drift no more than one major version, e.g. don't be older than 23 once 24 is available.

## Sams question

A DAG of three nodes: fetch -> envoy -> deliver. Submit host and cluster are far apart, and we need to propagate large quantities of data from one node to the next. How do we make this transfer quickly (i.e. without going through the submit host) without knowing the data's location at submit time?

krowe: Why do this as a dag? Why not make it one job instead of a dag? Collapsing the DAG into just one job has the advantage that it can use the local condor scratch area and can easily restart if the job fails without need for cleaning up anything. And of course making it one job means all the steps know where the data is.

Greg: condor\_chirp **condor\_chirp\_set\_job\_attr attributeName 'Value'** You could do something like

**condor\_chirp\_set\_job\_attr DataLocation ""(path)/to/something""**

or

**condor\_chirp put\_file local remote**

Each DAG has a prescript that runs before the dag nodes.

Another idea is to define the directory before submitting the job (e.g. /lustre/naasc/.../jobid)

## Condor history for crashed node

We have nodes crashing sometimes. 1. should HTCondor recover from a crashed node? Will the jobs be restarted somewhere else? 2. How can I see what jobs were running on a node when it crashed?

How about this

```
condor_history -name mcilroy -const "stringListMember('alias=nmpost091.aoc.nrao.edu/', StarterIpAddr, 'I&I') == true"
```

ANSWER: There is a global event log but it has to be enabled and isn't in our case **EVENT\_LOG = \$(LOG)/Event.log**

ANSWER: show jobs that have restarted **condor\_q -name mcilroy -allusers -const 'NumShadowStarts > 1'**

## STARTD\_ATTRS in glidein nodes

We add the following line to /etc/condor/condor\_config on all our Slurm nodes so that if they get called as a glidein node, they can set some special glidein settings.

```
LOCAL_CONFIG_FILE = /var/run/condor/condor_config.local
```

Our /etc/condor/config.d/99-nrao file effectively sets the following

```
STARTD_ATTRS = PoolName NRAO_TRANSFER_HOST HASLUSTRE BATCH
```

Our `/var/run/condor/condor_config.local`, which is run by glidein nodes, sets the following

```
STARTD_ATTRS = $(STARTD_ATTRS) NRAOGLIDEIN
```

The problem is glidein nodes don't get all the `STARTD_ATTRS` set by 99-nrao. They just get `NRAOGLIDEIN`. It is like condor-master reads 99-nrao to set its `STARTD_ATTRS`. Then it read `condor_config.local` to set its `STARTD_ATTRS` again but without accessing `$(STARTD_ATTRS)`.

ANSWER: The last line in `/var/run/condor/condor_config.local` is re-writing `STARTD_ATTRS`. It should have `$(STARTD_ATTRS)` appended

```
STARTD_ATTRS = NRAOGLIDEIN
```

Output to two places

Some of our pipeline jobs don't set `shoud_transfer_files=YES` because they need to transfer some output to an area for Analysts to look at and a some other output (may be a subset) to a different area for the User to look at. Is there a condor way to do this? `transfer_output_remaps`?

ANSWER: Greg doesn't think there is a Condor way to do this. Could make a copy of the subset and use `transfer_output_remaps` on the copy but that is a bit of a hack.

Pelican?

Felipe is playing with it and we will probably want it at NRAO.

ANSWER: Greg will ask around.

RHEL8 Crashing

We have had many NMT VLASS nodes crash since we upgraded to RHEL8. I think the nodes were busy when they crashed. So I changed our `SLOT_TYPE_1` from 100% to 95%. Is this a good idea?

ANSWER: try using `RESERVED_MEMORY=4096` (units are in Megabytes) instead of `SLOT_TYPE_1=95%` and put `SLOT_TYPE_1=100%` again.

No labels